



US012067260B2

(12) **United States Patent**
Hayes et al.

(10) **Patent No.:** **US 12,067,260 B2**
(45) **Date of Patent:** ***Aug. 20, 2024**

(54) **TRANSACTION PROCESSING WITH
DIFFERING CAPACITY STORAGE**

(71) Applicant: **Pure Storage, Inc.**, Mountain View, CA (US)

(72) Inventors: **John Hayes**, Mountain View, CA (US);
Robert Lee, San Carlos, CA (US); **Igor Ostrovsky**, Mountain View, CA (US);
Peter Vajgel, Menlo Park, CA (US)

(73) Assignee: **PURE STORAGE, INC.**, Santa Clara, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **17/548,112**

(22) Filed: **Dec. 10, 2021**

(65) **Prior Publication Data**

US 2022/0100399 A1 Mar. 31, 2022

Related U.S. Application Data

(63) Continuation of application No. 16/826,740, filed on Mar. 23, 2020, now Pat. No. 11,204,701, which is a (Continued)

(51) **Int. Cl.**
G06F 3/06 (2006.01)

(52) **U.S. Cl.**
CPC **G06F 3/0622** (2013.01); **G06F 3/0637** (2013.01); **G06F 3/0659** (2013.01); **G06F 3/067** (2013.01); **G06F 3/0688** (2013.01)

(58) **Field of Classification Search**

CPC **G06F 3/0622**; **G06F 3/0637**; **G06F 3/0659**;
G06F 3/067; **G06F 3/0688**

See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,390,327 A 2/1995 Lubbers et al.
5,450,581 A 9/1995 Bergen et al.

(Continued)

FOREIGN PATENT DOCUMENTS

EP 2164006 3/2010
EP 2256621 12/2010

(Continued)

OTHER PUBLICATIONS

Hwang, Kai, et al. "RAID-x: A New Distributed Disk Array for I/O-Centric Cluster Computing," HPDC '00 Proceedings of the 9th IEEE International Symposium on High Performance Distributed Computing, IEEE, 2000, pp. 279-286.

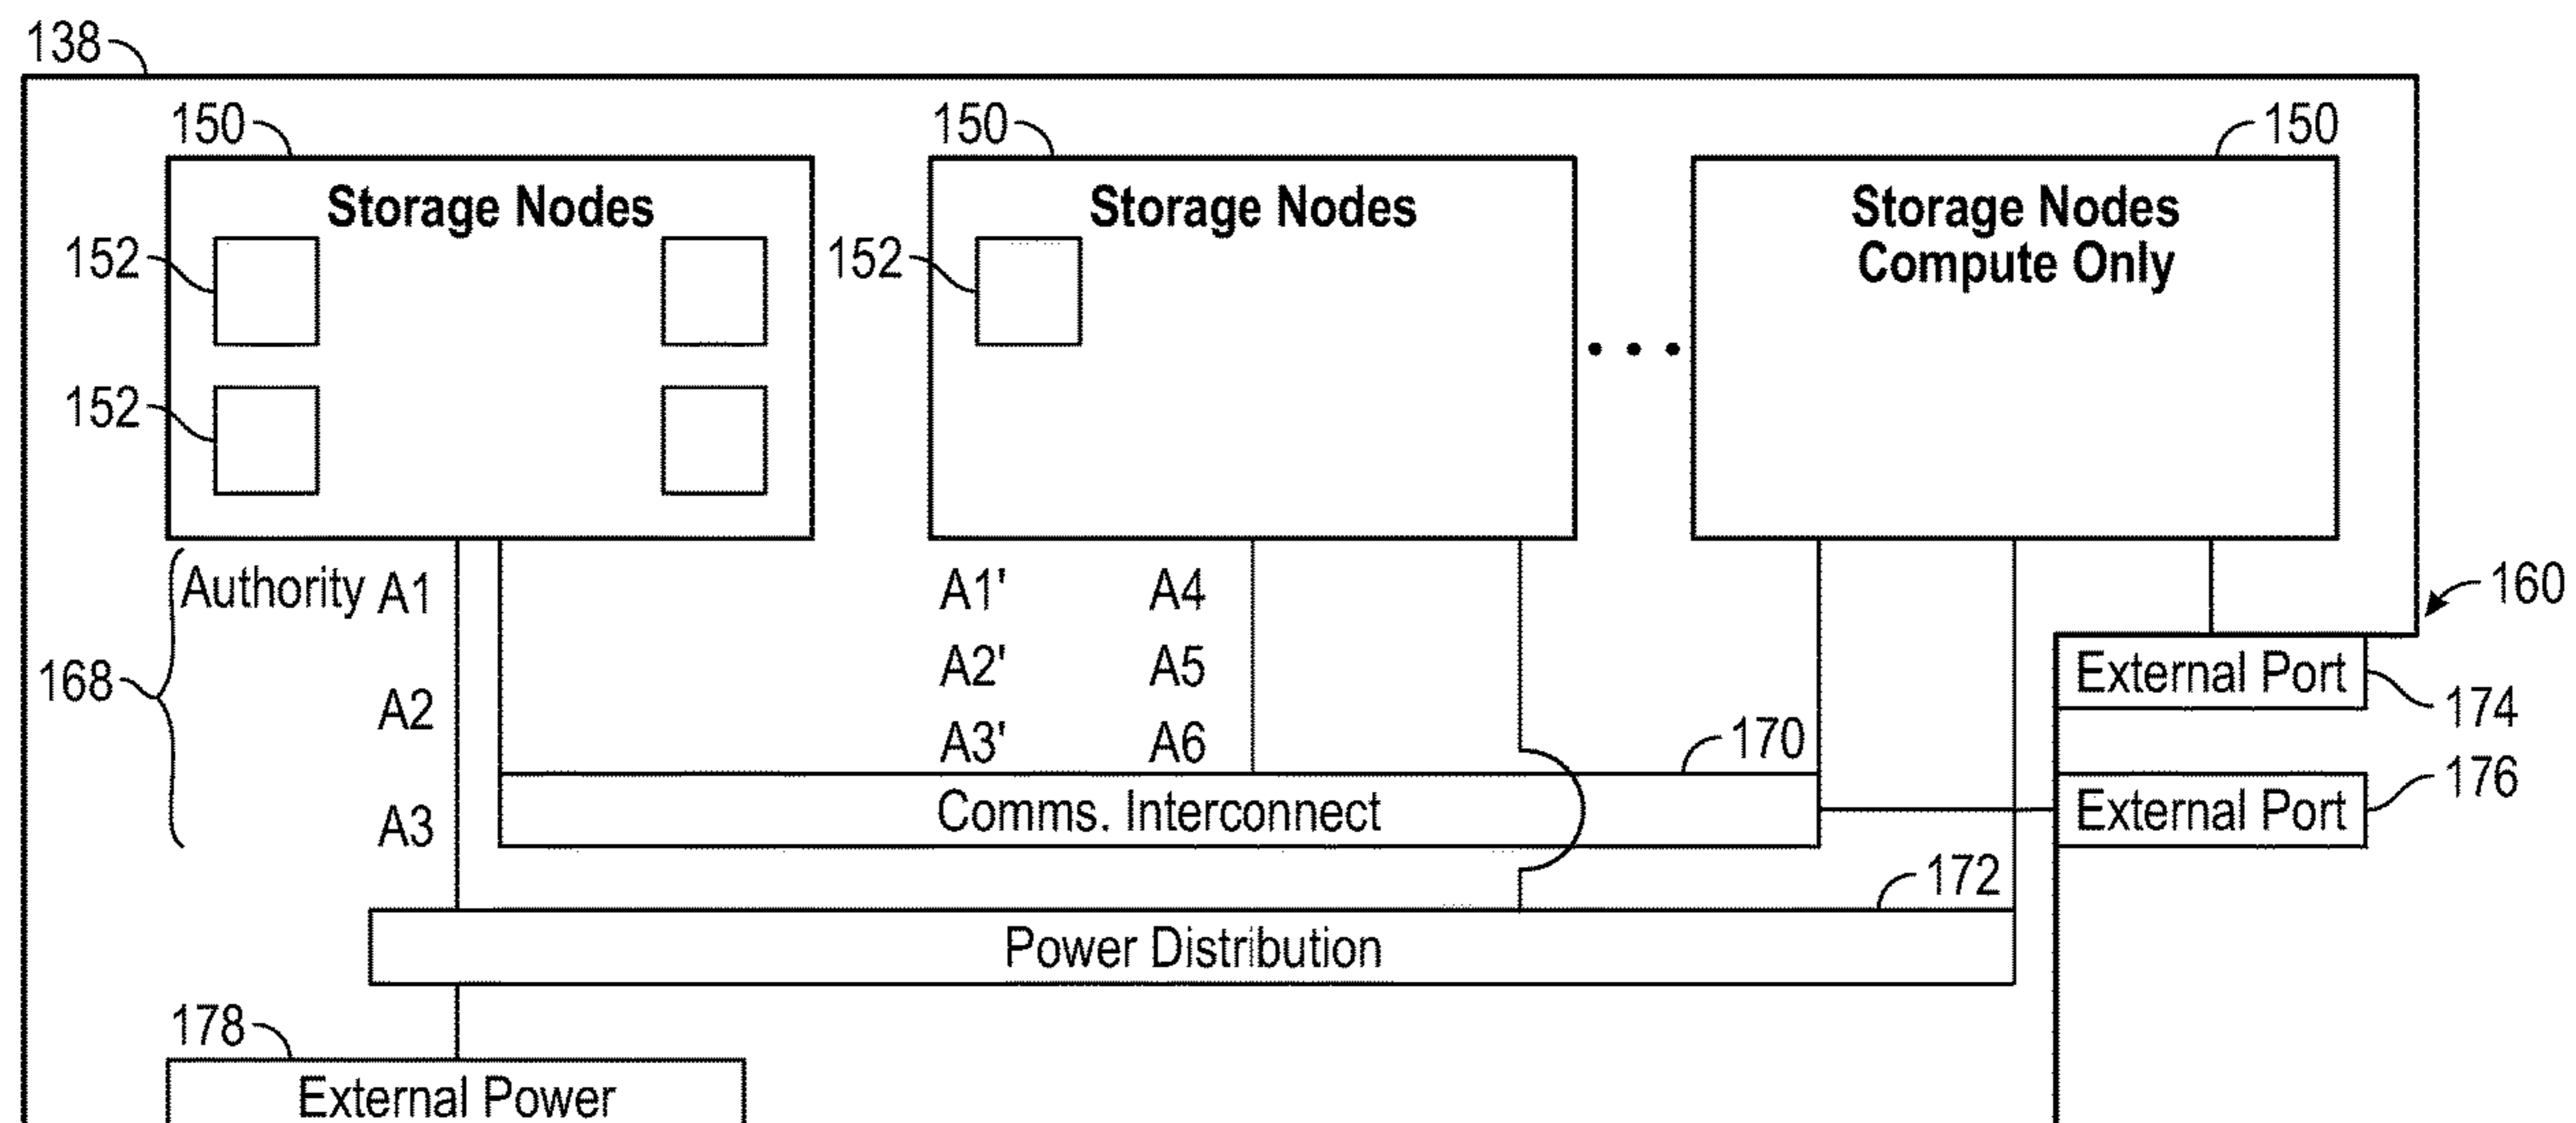
(Continued)

Primary Examiner — Jae U Yu

(57) **ABSTRACT**

A method of processing transactions associated with a command in a storage system is provided. The method includes receiving, at a first authority of the storage system, a command relating to user data. The method includes sending a transaction of the command, from the first authority to a second authority of the storage system, wherein a token accompanies the transaction and writing data in accordance with the transaction as permitted by the token into a partition that is allocated to the second authority in a storage device of the storage system.

20 Claims, 7 Drawing Sheets



Related U.S. Application Data

continuation of application No. 16/017,968, filed on Jun. 25, 2018, now Pat. No. 10,599,348, which is a continuation of application No. 14/978,073, filed on Dec. 22, 2015, now Pat. No. 10,007,457.

(56)

References Cited

U.S. PATENT DOCUMENTS

5,479,653	A	12/1995	Jones	7,870,105	B2	1/2011	Arakawa et al.
5,488,731	A	1/1996	Mendelsohn	7,873,878	B2	1/2011	Belluomini et al.
5,504,858	A	4/1996	Ellis et al.	7,885,938	B1	2/2011	Greene et al.
5,564,113	A	10/1996	Bergen et al.	7,886,111	B2	2/2011	Klemm et al.
5,574,882	A	11/1996	Menon et al.	7,908,448	B1	3/2011	Chatterjee et al.
5,649,093	A	7/1997	Hanko et al.	7,916,538	B2	3/2011	Jeon et al.
5,883,909	A	3/1999	Dekoning et al.	7,921,268	B2	4/2011	Jakob
6,000,010	A	12/1999	Legg	7,930,499	B2	4/2011	Duchesne
6,260,156	B1	7/2001	Garvin et al.	7,941,697	B2	5/2011	Mathew et al.
6,269,453	B1	7/2001	Krantz	7,958,303	B2	6/2011	Shuster
6,275,898	B1	8/2001	DeKoning	7,971,129	B2	6/2011	Watson
6,453,428	B1	9/2002	Stephenson	7,984,016	B2	7/2011	Kisley
6,523,087	B2	2/2003	Busser	7,991,822	B2	8/2011	Bish et al.
6,535,417	B2	3/2003	Tsuda	8,006,126	B2	8/2011	Deenadhayalan et al.
6,643,748	B1	11/2003	Wieland	8,010,485	B1	8/2011	Chatterjee et al.
6,725,392	B1	4/2004	Frey et al.	8,010,829	B1	8/2011	Chatterjee et al.
6,763,455	B2	7/2004	Hall	8,020,047	B2	9/2011	Courtney
6,836,816	B2	12/2004	Kendall	8,046,548	B1	10/2011	Chatterjee et al.
6,985,995	B2	1/2006	Holland et al.	8,051,361	B2	11/2011	Sim-Tang et al.
7,032,125	B2	4/2006	Holt et al.	8,051,362	B2	11/2011	Li et al.
7,047,358	B2	5/2006	Lee et al.	8,074,038	B2	12/2011	Lionetti et al.
7,051,155	B2	5/2006	Talagala et al.	8,082,393	B2	12/2011	Galloway et al.
7,055,058	B2	5/2006	Lee et al.	8,086,603	B2	12/2011	Nasre et al.
7,065,617	B2	6/2006	Wang	8,086,634	B2	12/2011	Mimatsu
7,069,383	B2	6/2006	Yamamoto et al.	8,086,911	B1	12/2011	Taylor
7,076,606	B2	7/2006	Orsley	8,090,837	B2	1/2012	Shin et al.
7,107,480	B1	9/2006	Moshayedi et al.	8,108,502	B2	1/2012	Tabbara et al.
7,159,150	B2	1/2007	Kenchamma-Hosekote et al.	8,117,388	B2	2/2012	Jernigan, IV
7,162,575	B2	1/2007	Dalal et al.	8,117,521	B2	2/2012	Yang et al.
7,164,608	B2	1/2007	Lee	8,140,821	B1	3/2012	Raizen et al.
7,188,270	B1	3/2007	Nanda et al.	8,145,838	B1	3/2012	Miller et al.
7,334,156	B2	2/2008	Land et al.	8,145,840	B2	3/2012	Koul et al.
7,370,220	B1	5/2008	Nguyen et al.	8,175,012	B2	5/2012	Haratsch et al.
7,386,666	B1	6/2008	Beauchamp et al.	8,176,360	B2	5/2012	Frost et al.
7,398,285	B2	7/2008	Kisley	8,176,405	B2	5/2012	Hafner et al.
7,424,498	B1	9/2008	Patterson	8,180,855	B2	5/2012	Aiello et al.
7,424,592	B1	9/2008	Karr	8,200,922	B2	6/2012	McKean et al.
7,444,532	B2	10/2008	Masuyama et al.	8,209,469	B2	6/2012	Carpenter et al.
7,480,658	B2	1/2009	Heinla et al.	8,225,006	B1	7/2012	Karamcheti
7,484,056	B2	1/2009	Madnani et al.	8,239,618	B2	8/2012	Kotzur et al.
7,484,057	B1	1/2009	Madnani et al.	8,244,999	B1	8/2012	Chatterjee et al.
7,484,059	B1	1/2009	Ofer et al.	8,261,016	B1	9/2012	Goel
7,536,506	B2	5/2009	Ashmore et al.	8,271,455	B2	9/2012	Kesselman
7,558,859	B2	7/2009	Kasiolas	8,285,686	B2	10/2012	Kesselman
7,565,446	B2	7/2009	Talagala et al.	8,305,811	B2	11/2012	Jeon
7,613,947	B1	11/2009	Coatney	8,315,999	B2	11/2012	Chatley et al.
7,634,617	B2	12/2009	Misra	8,327,080	B1	12/2012	Der
7,634,618	B2	12/2009	Misra	8,335,769	B2	12/2012	Kesselman
7,681,104	B1	3/2010	Sim-Tang et al.	8,341,118	B2	12/2012	Drobychev et al.
7,681,105	B1	3/2010	Sim-Tang et al.	8,351,290	B1	1/2013	Huang et al.
7,681,109	B2	3/2010	Yang et al.	8,364,920	B1	1/2013	Parkison et al.
7,730,257	B2	6/2010	Franklin	8,365,041	B2	1/2013	Chu et al.
7,730,258	B1	6/2010	Smith	8,375,146	B2	2/2013	Sinclair
7,730,274	B1	6/2010	Usgaonkar	8,397,016	B2	3/2013	Talagala et al.
7,743,276	B2	6/2010	Jacobsen et al.	8,402,152	B2	3/2013	Duran
7,752,489	B2	7/2010	Deenadhayalan et al.	8,412,880	B2	4/2013	Leibowitz et al.
7,757,038	B2	7/2010	Kitahara	8,423,739	B2	4/2013	Ash et al.
7,757,059	B1	7/2010	Ofer et al.	8,429,436	B2	4/2013	Filingim et al.
7,778,960	B1	8/2010	Chatterjee et al.	8,452,928	B1	5/2013	Ofer et al.
7,783,955	B2	8/2010	Haratsch et al.	8,473,698	B2	6/2013	Lionetti et al.
7,814,272	B2	10/2010	Barrall et al.	8,473,778	B2	6/2013	Simitci
7,814,273	B2	10/2010	Barrall	8,473,815	B2	6/2013	Yu et al.
7,818,531	B2	10/2010	Barrall	8,479,037	B1	7/2013	Chatterjee et al.
7,827,351	B2	11/2010	Suetsugu et al.	8,484,414	B2	7/2013	Sugimoto et al.
7,827,439	B2	11/2010	Matthew et al.	8,498,967	B1	7/2013	Chatterjee et al.
7,831,768	B2	11/2010	Ananthamurthy et al.	8,522,073	B2	8/2013	Cohen
7,856,583	B1	12/2010	Smith	8,533,408	B1	9/2013	Madnani et al.
				8,533,527	B2	9/2013	Daikokuya et al.
				8,539,177	B1	9/2013	Ofer et al.
				8,544,029	B2	9/2013	Bakke et al.
				8,549,224	B1	10/2013	Zeryck et al.
				8,583,861	B1	11/2013	Ofer et al.
				8,589,625	B2	11/2013	Colgrove et al.
				8,595,455	B2	11/2013	Chatterjee et al.
				8,615,599	B1	12/2013	Takefiman et al.
				8,627,136	B2	1/2014	Shankar et al.
				8,627,138	B1	1/2014	Clark
				8,639,669	B1	1/2014	Douglis et al.
				8,639,863	B1	1/2014	Kanapathippillai et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

8,640,000 B1	1/2014	Cypher	9,123,422 B2	9/2015	Sharon et al.
8,650,343 B1	2/2014	Kanapathippillai et al.	9,124,300 B2	9/2015	Olbrich et al.
8,660,131 B2	2/2014	Vermunt et al.	9,134,908 B2	9/2015	Horn et al.
8,661,218 B1	2/2014	Piszczek et al.	9,153,337 B2	10/2015	Sutardja
8,671,072 B1	3/2014	Shah et al.	9,158,472 B2	10/2015	Kesselman et al.
8,671,445 B1	3/2014	Wang et al.	9,159,422 B1	10/2015	Lee et al.
8,689,042 B1	4/2014	Kanapathippillai et al.	9,164,891 B2	10/2015	Karamcheti et al.
8,700,875 B1	4/2014	Barron et al.	9,183,136 B2	11/2015	Kawamura et al.
8,706,694 B2	4/2014	Chatterjee et al.	9,189,650 B2	11/2015	Jaye et al.
8,706,914 B2	4/2014	Duchesneau	9,201,733 B2	12/2015	Verma
8,706,932 B1	4/2014	Kanapathippillai et al.	9,207,876 B2	12/2015	Shu et al.
8,712,963 B1	4/2014	Douglis et al.	9,229,656 B1	1/2016	Contreras et al.
8,713,405 B2	4/2014	Healey et al.	9,229,810 B2	1/2016	He et al.
8,719,621 B1	5/2014	Karmarkar	9,235,475 B1	1/2016	Shilane et al.
8,725,730 B2	5/2014	Keeton et al.	9,244,626 B2	1/2016	Shah et al.
8,751,859 B2	6/2014	Becker-szendy et al.	9,250,999 B1	2/2016	Barroso
8,756,387 B2	6/2014	Frost et al.	9,251,066 B2	2/2016	Colgrove et al.
8,762,793 B2	6/2014	Grube et al.	9,268,648 B1	2/2016	Barash et al.
8,838,541 B2	6/2014	Camble et al.	9,268,806 B1	2/2016	Kesselman et al.
8,769,232 B2	7/2014	Suryabudi et al.	9,286,002 B1	3/2016	Karamcheti et al.
8,775,858 B2	7/2014	Gower et al.	9,292,214 B2	3/2016	Kalos et al.
8,775,868 B2	7/2014	Colgrove et al.	9,298,760 B1	3/2016	Li et al.
8,788,913 B1	7/2014	Xin et al.	9,304,908 B1	4/2016	Karamcheti et al.
8,793,447 B2	7/2014	Usgaonkar et al.	9,311,969 B2	4/2016	Murin
8,799,746 B2	8/2014	Baker et al.	9,311,970 B2	4/2016	Sharon et al.
8,819,311 B2	8/2014	Liao	9,323,663 B2	4/2016	Karamcheti et al.
8,819,383 B1	8/2014	Jobanputra et al.	9,323,667 B2	4/2016	Bennett
8,824,261 B1	9/2014	Miller et al.	9,323,681 B2	4/2016	Apostolides et al.
8,832,528 B2	9/2014	Thatcher et al.	9,335,942 B2	5/2016	Kumar et al.
8,838,892 B2	9/2014	Li	9,348,538 B2	5/2016	Mallaiah et al.
8,843,700 B1	9/2014	Salessi et al.	9,355,022 B2	5/2016	Ravimohan et al.
8,850,108 B1	9/2014	Hayes et al.	9,384,082 B1	7/2016	Lee et al.
8,850,288 B1	9/2014	Lazier et al.	9,384,252 B2	7/2016	Akirav et al.
8,856,593 B2	10/2014	Eckhardt et al.	9,389,958 B2	7/2016	Sundaram et al.
8,856,619 B1	10/2014	Cypher	9,390,019 B2	7/2016	Patterson et al.
8,862,617 B2	10/2014	Kesselman	9,396,202 B1	7/2016	Drobychev et al.
8,862,847 B2	10/2014	Feng et al.	9,400,828 B2	7/2016	Kesselman et al.
8,862,928 B2	10/2014	Xavier et al.	9,405,478 B2	8/2016	Koseki et al.
8,868,825 B1	10/2014	Hayes	9,411,685 B2	8/2016	Lee
8,874,836 B1	10/2014	Hayes	9,417,960 B2	8/2016	Klein
8,880,793 B2	11/2014	Nagineni	9,417,963 B2	8/2016	He et al.
8,880,825 B2	11/2014	Lionetti et al.	9,430,250 B2	8/2016	Hamid et al.
8,886,778 B2	11/2014	Nedved et al.	9,430,542 B2	8/2016	Akirav et al.
8,898,383 B2	11/2014	Yamamoto et al.	9,432,541 B2	8/2016	Ishida
8,898,388 B1	11/2014	Kimmel	9,454,434 B2	9/2016	Sundaram et al.
8,904,231 B2	12/2014	Coatney et al.	9,471,579 B1	10/2016	Natanzon
8,918,478 B2	12/2014	Ozzie et al.	9,477,554 B2	10/2016	Chamness et al.
8,930,307 B2	1/2015	Colgrove et al.	9,477,632 B2	10/2016	Du
8,930,633 B2	1/2015	Amit et al.	9,501,398 B2	11/2016	George et al.
8,943,357 B2	1/2015	Atzmony	9,525,737 B2	12/2016	Friedman
8,949,502 B2	2/2015	McKnight et al.	9,529,542 B2	12/2016	Friedman et al.
8,959,110 B2	2/2015	Smith et al.	9,535,631 B2	1/2017	Fu et al.
8,959,388 B1	2/2015	Kuang et al.	9,552,248 B2	1/2017	Miller et al.
8,972,478 B1	3/2015	Storer et al.	9,552,291 B2	1/2017	Munetoh et al.
8,972,779 B2	3/2015	Lee et al.	9,552,299 B2	1/2017	Stalzer
8,977,597 B2	3/2015	Ganesh et al.	9,563,517 B1	2/2017	Natanzon et al.
8,996,828 B2	3/2015	Kalos et al.	9,588,698 B1	3/2017	Karamcheti et al.
9,003,144 B1	4/2015	Hayes et al.	9,588,712 B2	3/2017	Kalos et al.
9,009,724 B2	4/2015	Gold et al.	9,594,652 B1	3/2017	Sathiamoorthy et al.
9,021,053 B2	4/2015	Bernbo et al.	9,600,193 B2	3/2017	Ahrens et al.
9,021,215 B2	4/2015	Meir et al.	9,619,321 B1	4/2017	Sharon et al.
9,021,297 B1	4/2015	Hayes et al.	9,619,430 B2	4/2017	Kannan et al.
9,025,393 B2	5/2015	Wu	9,645,754 B2	5/2017	Li et al.
9,043,372 B2	5/2015	Makkar et al.	9,667,720 B1	5/2017	Bent et al.
9,047,214 B1	6/2015	Sharon et al.	9,710,535 B2	7/2017	Aizman et al.
9,053,808 B2	6/2015	Sprouse	9,733,840 B2	8/2017	Karamcheti et al.
9,058,155 B2	6/2015	Cepulis et al.	9,734,225 B2	8/2017	Akirav et al.
9,063,895 B1	6/2015	Madnani et al.	9,740,403 B2	8/2017	Storer et al.
9,063,896 B1	6/2015	Madnani et al.	9,740,700 B1	8/2017	Chopra et al.
9,098,211 B1	8/2015	Madnani et al.	9,740,762 B2	8/2017	Horowitz et al.
9,110,898 B1	8/2015	Chamness et al.	9,747,319 B2	8/2017	Bestler et al.
9,110,964 B1	8/2015	Shilane et al.	9,747,320 B2	8/2017	Kesselman
9,116,819 B2	8/2015	Cope et al.	9,767,130 B2	9/2017	Bestler et al.
9,117,536 B2	8/2015	Yoon	9,781,227 B2	10/2017	Friedman et al.
9,122,401 B2	9/2015	Zaltsman et al.	9,785,498 B2	10/2017	Misra et al.
			9,798,486 B1	10/2017	Singh
			9,804,925 B1	10/2017	Carmi et al.
			9,811,285 B1	11/2017	Karamcheti et al.
			9,811,546 B1	11/2017	Bent et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

9,818,478 B2	11/2017	Chung	2010/0162076 A1	6/2010	Sim-Tang et al.
9,829,066 B2	11/2017	Thomas	2010/0169707 A1	7/2010	Mathew et al.
9,836,245 B2	12/2017	Hayes et al.	2010/0174576 A1	7/2010	Naylor
9,891,854 B2	2/2018	Munetoh et al.	2010/0268908 A1	10/2010	Ouyang et al.
9,891,860 B1	2/2018	Delgado et al.	2011/0040925 A1	2/2011	Frost et al.
9,892,005 B2	2/2018	Kedem et al.	2011/0060927 A1	3/2011	Fillingim et al.
9,892,186 B2	2/2018	Akirav et al.	2011/0119462 A1	5/2011	Leach et al.
9,904,589 B1	2/2018	Donlan et al.	2011/0202721 A1	8/2011	Popovski et al.
9,904,717 B2	2/2018	Anglin et al.	2011/0219170 A1	9/2011	Frost et al.
9,952,809 B2	2/2018	Shah	2011/0238625 A1	9/2011	Hamaguchi et al.
9,910,748 B2	3/2018	Pan	2011/0264843 A1	10/2011	Haines et al.
9,910,904 B2	3/2018	Anglin et al.	2011/0302369 A1	12/2011	Goto et al.
9,934,237 B1	4/2018	Shilane et al.	2012/0011398 A1	1/2012	Eckhardt
9,940,065 B2	4/2018	Kalos et al.	2012/0079318 A1	3/2012	Colgrove et al.
9,946,604 B1	4/2018	Glass	2012/0089567 A1	4/2012	Takahashi et al.
9,959,167 B1	5/2018	Donlan et al.	2012/0110249 A1	5/2012	Jeong et al.
9,965,539 B2	5/2018	D'halluin et al.	2012/0131253 A1	5/2012	McKnight
9,998,539 B1	6/2018	Brock et al.	2012/0158923 A1	6/2012	Mohamed et al.
10,007,457 B2	6/2018	Hayes et al.	2012/0191900 A1	7/2012	Kunimatsu et al.
10,013,177 B2	7/2018	Liu et al.	2012/0198152 A1	8/2012	Terry et al.
10,013,311 B2	7/2018	Sundaram et al.	2012/0198261 A1	8/2012	Brown et al.
10,019,314 B2	7/2018	Litsyn et al.	2012/0209943 A1	8/2012	Jung
10,019,317 B2	7/2018	Usvyatsky et al.	2012/0226772 A1	9/2012	Grube et al.
10,031,703 B1	7/2018	Natanzon et al.	2012/0226934 A1	9/2012	Rao
10,061,512 B2	8/2018	Chu et al.	2012/0246435 A1	9/2012	Meir et al.
10,073,626 B2	9/2018	Karamcheti et al.	2012/0260055 A1	10/2012	Murase
10,082,985 B2	9/2018	Hayes et al.	2012/0311557 A1	12/2012	Resch
10,089,012 B1	10/2018	Chen et al.	2013/0022201 A1	1/2013	Glew et al.
10,089,174 B2	10/2018	Lin	2013/0036314 A1	2/2013	Glew et al.
10,089,176 B1	10/2018	Donlan et al.	2013/0042056 A1	2/2013	Shats
10,108,819 B1	10/2018	Donlan et al.	2013/0060884 A1	3/2013	Bernbo et al.
10,146,787 B2	12/2018	Bashyam et al.	2013/0067188 A1	3/2013	Mehra et al.
10,152,268 B1	12/2018	Chakraborty et al.	2013/0073894 A1	3/2013	Xavier et al.
10,157,098 B2	12/2018	Chung et al.	2013/0124776 A1	5/2013	Hallak et al.
10,162,704 B1	12/2018	Kirschner et al.	2013/0132800 A1	5/2013	Healy et al.
10,180,875 B2	1/2019	Northcott	2013/0151653 A1	6/2013	Sawiki
10,185,730 B2	1/2019	Bestler et al.	2013/0151771 A1	6/2013	Tsukahara et al.
10,235,065 B1	3/2019	Miller et al.	2013/0173853 A1	7/2013	Ungureanu et al.
2002/0112178 A1*	8/2002	Scherr G06F 21/80 726/4	2013/0238554 A1	9/2013	Yucel et al.
2002/0144059 A1	10/2002	Kendall	2013/0339314 A1	12/2013	Carpenter et al.
2003/0105984 A1	6/2003	Masuyama et al.	2013/0339635 A1	12/2013	Amit et al.
2003/0110205 A1	6/2003	Johnson	2013/0339818 A1	12/2013	Baker et al.
2004/0161086 A1	8/2004	Buntin et al.	2014/0040535 A1	2/2014	Lee
2005/0001652 A1	1/2005	Malik et al.	2014/0040702 A1	2/2014	He et al.
2005/0076228 A1	4/2005	Davis et al.	2014/0047263 A1	2/2014	Coatney et al.
2005/0108730 A1	5/2005	Corrado et al.	2014/0047269 A1	2/2014	Kim
2005/0235132 A1	10/2005	Karr et al.	2014/0063721 A1	3/2014	Herman et al.
2005/0240726 A1	10/2005	Williams	2014/0064048 A1	3/2014	Cohen et al.
2005/0278460 A1	12/2005	Shin et al.	2014/0068224 A1	3/2014	Fan et al.
2005/0283649 A1	12/2005	Turner et al.	2014/0075252 A1	3/2014	Luo et al.
2006/0015683 A1	1/2006	Ashmore et al.	2014/0122510 A1	5/2014	Namkoong et al.
2006/0114930 A1	6/2006	Lucas et al.	2014/0136880 A1	5/2014	Shankar et al.
2006/0174157 A1	8/2006	Barrall et al.	2014/0181402 A1	6/2014	White
2006/0248294 A1	11/2006	Nedved et al.	2014/0229670 A1	8/2014	Madhusudana et al.
2007/0079068 A1	4/2007	Draggon	2014/0237164 A1	8/2014	Le et al.
2007/0214194 A1	9/2007	Reuter	2014/0279936 A1	9/2014	Bernbo et al.
2007/0214314 A1	9/2007	Reuter	2014/0280025 A1	9/2014	Eidson et al.
2007/0234016 A1	10/2007	Davis et al.	2014/0289588 A1	9/2014	Nagadomi et al.
2007/0268905 A1	11/2007	Baker et al.	2014/0330785 A1	11/2014	Isherwood et al.
2008/0080709 A1	4/2008	Michtchenko et al.	2014/0330921 A1	11/2014	Storm et al.
2008/0107274 A1	5/2008	Worthy	2014/0372838 A1	12/2014	Lou et al.
2008/0155191 A1	6/2008	Anderson et al.	2014/0380125 A1	12/2014	Calder et al.
2008/0288714 A1*	11/2008	Salomon G06F 8/61 711/E12.008	2014/0380126 A1	12/2014	Yekhanin et al.
2008/0295118 A1	11/2008	Liao	2015/0032720 A1	1/2015	James
2009/0077208 A1	3/2009	Nguyen et al.	2015/0039645 A1	2/2015	Lewis
2009/0138654 A1	5/2009	Sutardja	2015/0039849 A1	2/2015	Lewis
2009/0216910 A1	8/2009	Duchesneau	2015/0089283 A1	3/2015	Kermarrec et al.
2009/0216920 A1	8/2009	Lauterbach et al.	2015/0100746 A1	4/2015	Rychlik
2010/0017444 A1	1/2010	Chatterjee et al.	2015/0134824 A1	5/2015	Mickens
2010/0042636 A1	2/2010	Lu	2015/0153800 A1	6/2015	Lucas et al.
2010/0094806 A1	4/2010	Apostolides et al.	2015/0180714 A1	6/2015	Chunn
2010/0115070 A1	5/2010	Missimilly	2015/0280959 A1	10/2015	Vincent
2010/0125695 A1	5/2010	Wu et al.	2016/0246537 A1	2/2016	Kim
			2016/0191508 A1	6/2016	Bestler et al.
			2016/0378612 A1	12/2016	Hipsh et al.
			2017/0091236 A1	3/2017	Hayes et al.
			2017/0103092 A1	4/2017	Hu et al.
			2017/0103094 A1	4/2017	Hu et al.
			2017/0103098 A1	4/2017	Hu et al.

(56)

References Cited

U.S. PATENT DOCUMENTS

2017/0103116	A1	4/2017	Hu et al.
2017/0177236	A1	6/2017	Haratsch et al.
2018/0039442	A1	2/2018	Shadrin et al.
2018/0081958	A1	3/2018	Akirav et al.
2018/0101441	A1	4/2018	Hyun et al.
2018/0101587	A1	4/2018	Anglin et al.
2018/0101588	A1	4/2018	Anglin et al.
2018/0217756	A1	8/2018	Liu et al.
2018/0307560	A1	10/2018	Vishnumolakala et al.
2018/0321874	A1	11/2018	Li et al.
2019/0036703	A1	1/2019	Bestler

FOREIGN PATENT DOCUMENTS

WO	WO 02-13033	2/2002
WO	WO 2008103569	8/2008
WO	WO 2008157081	12/2008
WO	WO-2012166441	A1 12/2012
WO	WO 2013032825	7/2013

OTHER PUBLICATIONS

Schmid, Patrick: "RAID Scaling Charts, Part 3:4-128 kB Stripes Compared", Tom's Hardware, Nov. 27, 2007 (<http://www.tomshardware.com/reviews/RAID-SCALING-CHARTS.1735-4.html>), See pp. 1-2.

Storer, Mark W. et al., "Pergamum: Replacing Tape with Energy Efficient, Reliable, Disk Based Archival Storage," Fast '08: 6th USENIX Conference on File and Storage Technologies, San Jose, CA, Feb. 26-29, 2008 pp. 1-16.

Ju-Kyeong Kim et al., "Data Access Frequency based Data Replication Method using Erasure Codes in Cloud Storage System", Journal of the Institute of Electronics and Information Engineers, Feb. 2014, vol. 51, No. 2, pp. 85-91.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2015/018169, dated May 15, 2015.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2015/034302, dated Sep. 11, 2015.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2015/039135, dated Sep. 18, 2015.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2015/039136, dated Sep. 23, 2015.

International Search Report, PCT/US2015/039142, dated Sep. 24, 2015.

International Search Report, PCT/US2015/034291, dated Sep. 30, 2015.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2015/039137, dated Oct. 1, 2015.

International Search Report, PCT/US2015/044370, dated Dec. 15, 2015.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/031039, dated May 5, 2016.

International Search Report, PCT/US2016/014604, dated May 19, 2016.

International Search Report, PCT/US2016/014361, dated May 30, 2016.

International Search Report, PCT/US2016/014356, dated Jun. 28, 2016.

International Search Report, PCT/US2016/014357, dated Jun. 29, 2016.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/016504, dated Jul. 6, 2016.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/024391, dated Jul. 12, 2016.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/026529, dated Jul. 19, 2016.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/023485, dated Jul. 21, 2016.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/033306, dated Aug. 19, 2016.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/047808, dated Nov. 25, 2016.

Stalzer, Mark A., "FlashBlades: System Architecture and Applications," Proceedings of the 2nd Workshop on Architectures and Systems for Big Data, Association for Computing Machinery, New York, NY, 2012, pp. 10-14.

International Search Report and the Written Opinion of the International Searching Authority, PCT/US2016/042147, dated Nov. 30, 2016.

Extended European Search Report for European Application No. 16879995.5, mailed Apr. 26, 2019, 7 pages.

Extended European Search Report for European Application No. 23181330.4, mailed Sep. 21, 2023, 9 pages.

International Search Report and Written Opinion for International Application No. PCT/US2016/067877, mailed Mar. 31, 2017, 9 pages.

* cited by examiner

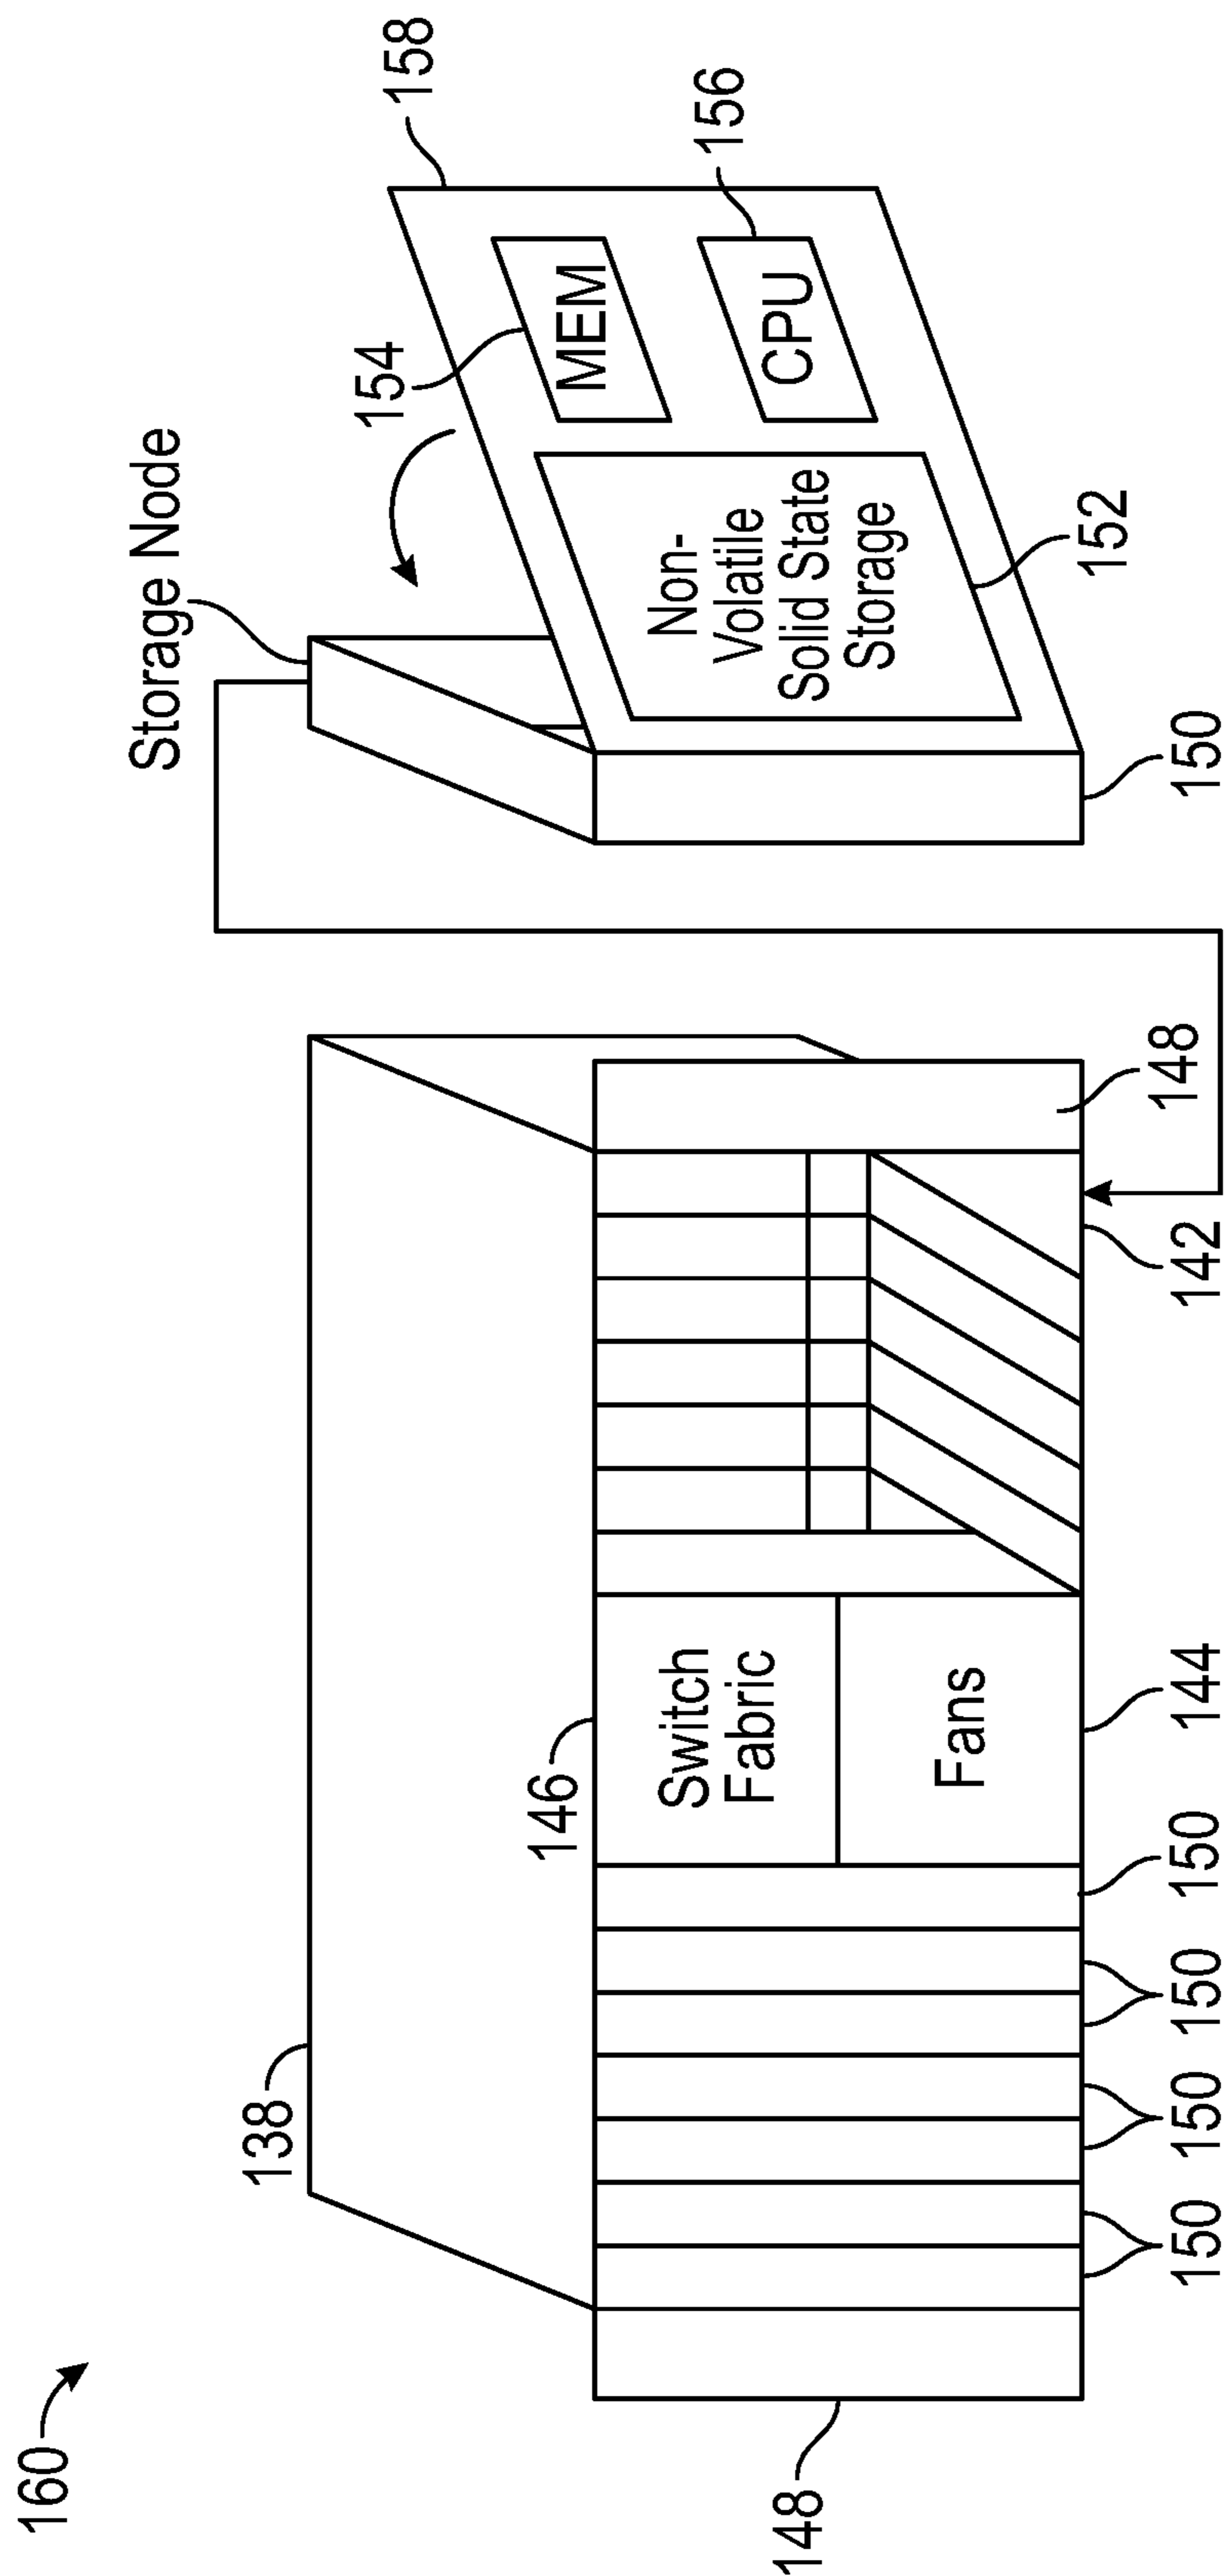


FIG. 1

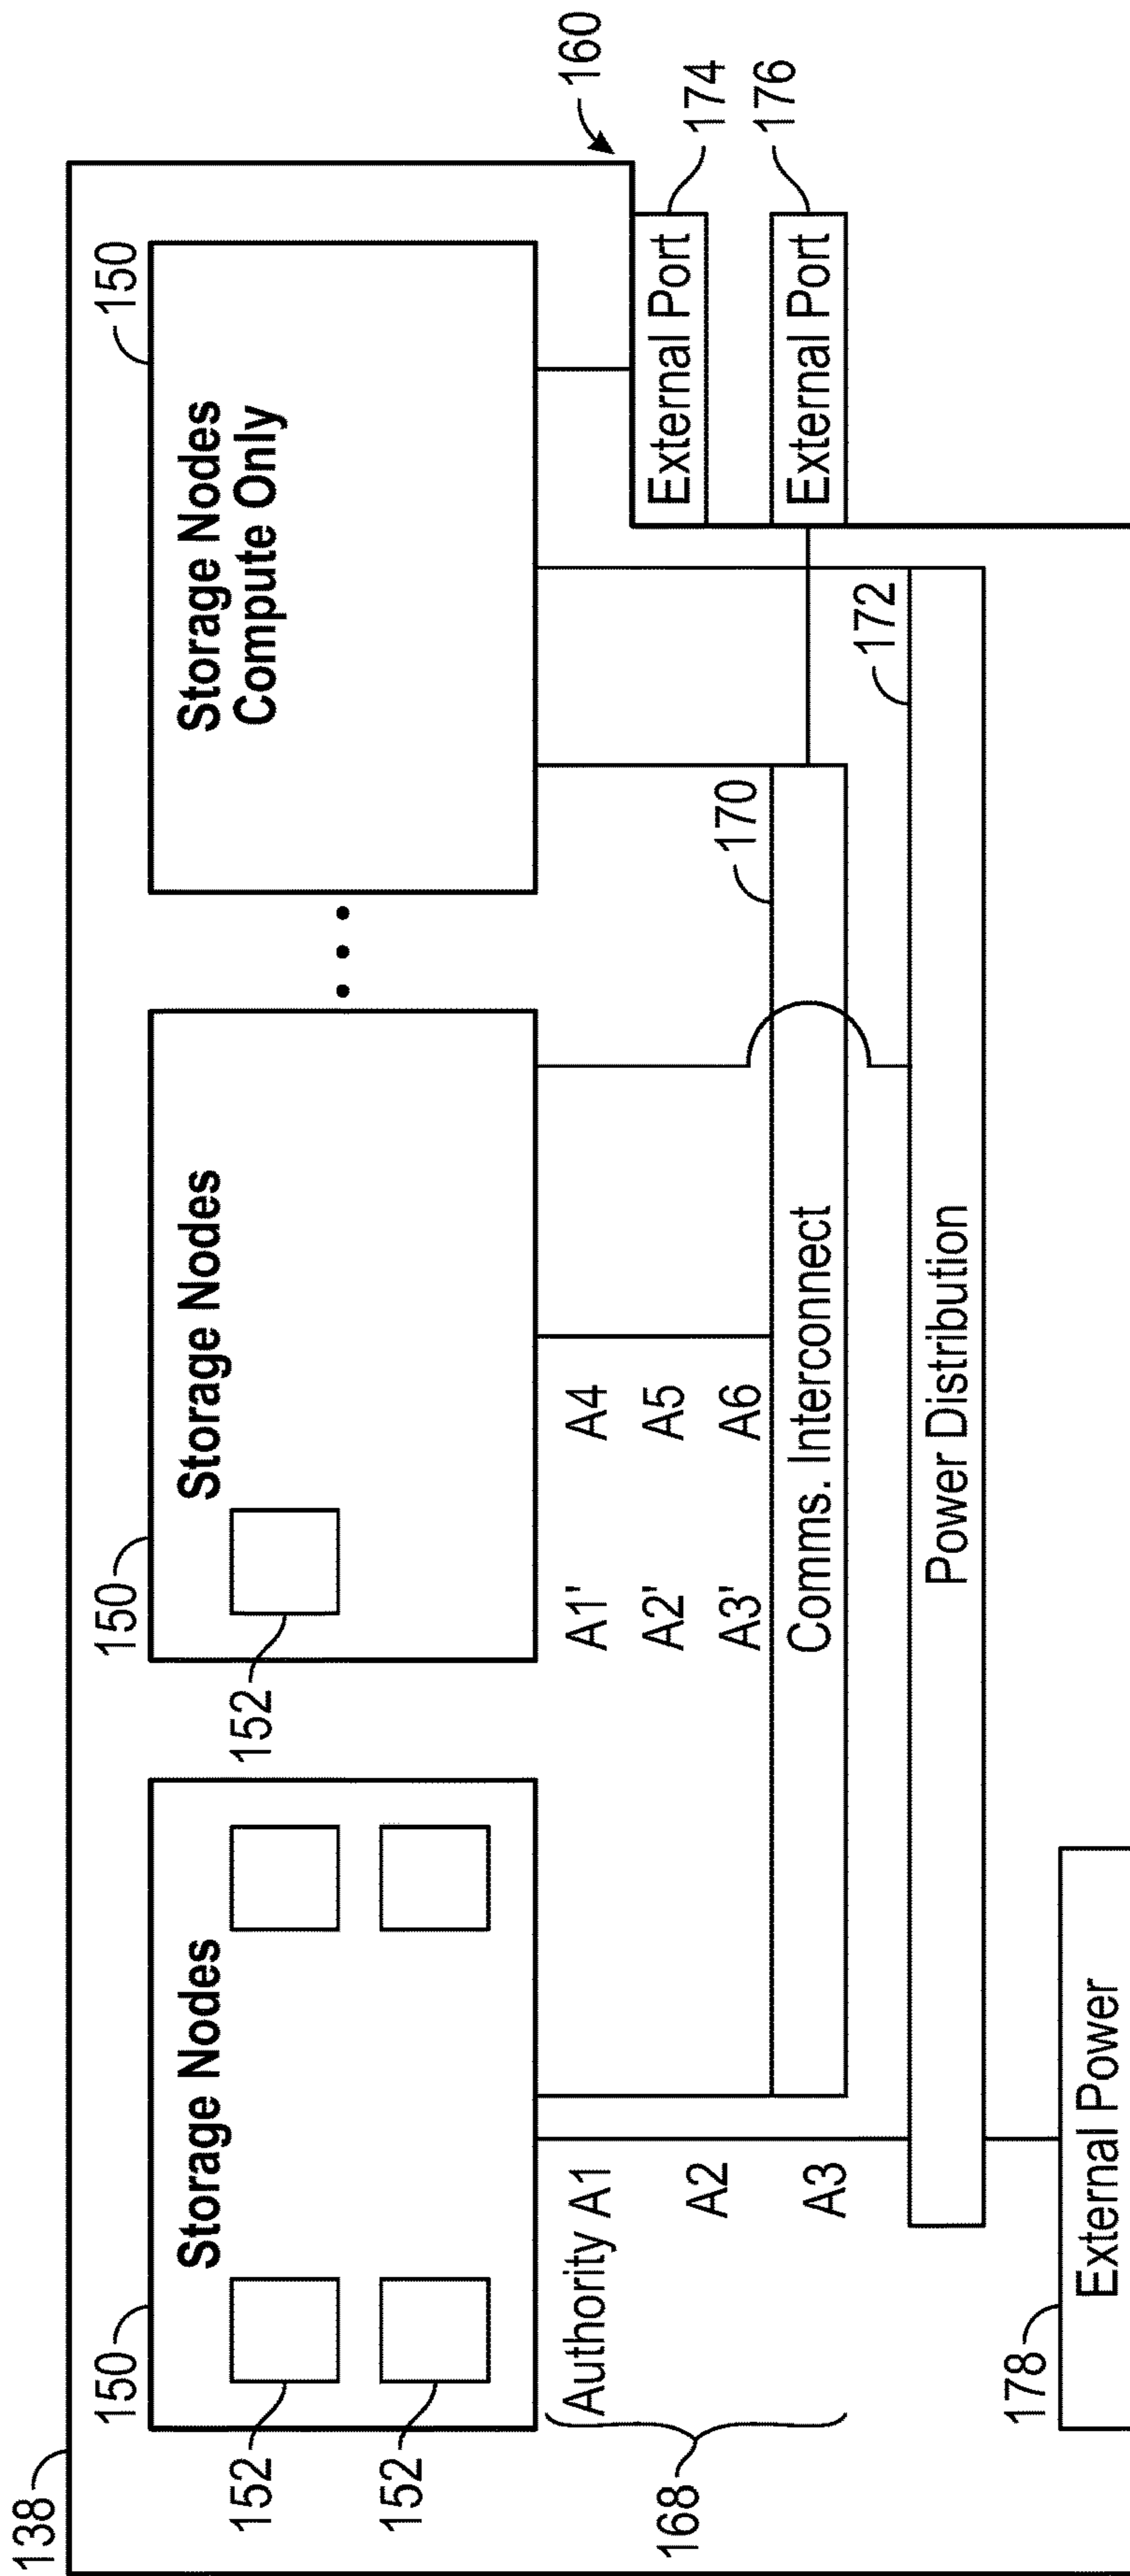


FIG. 2

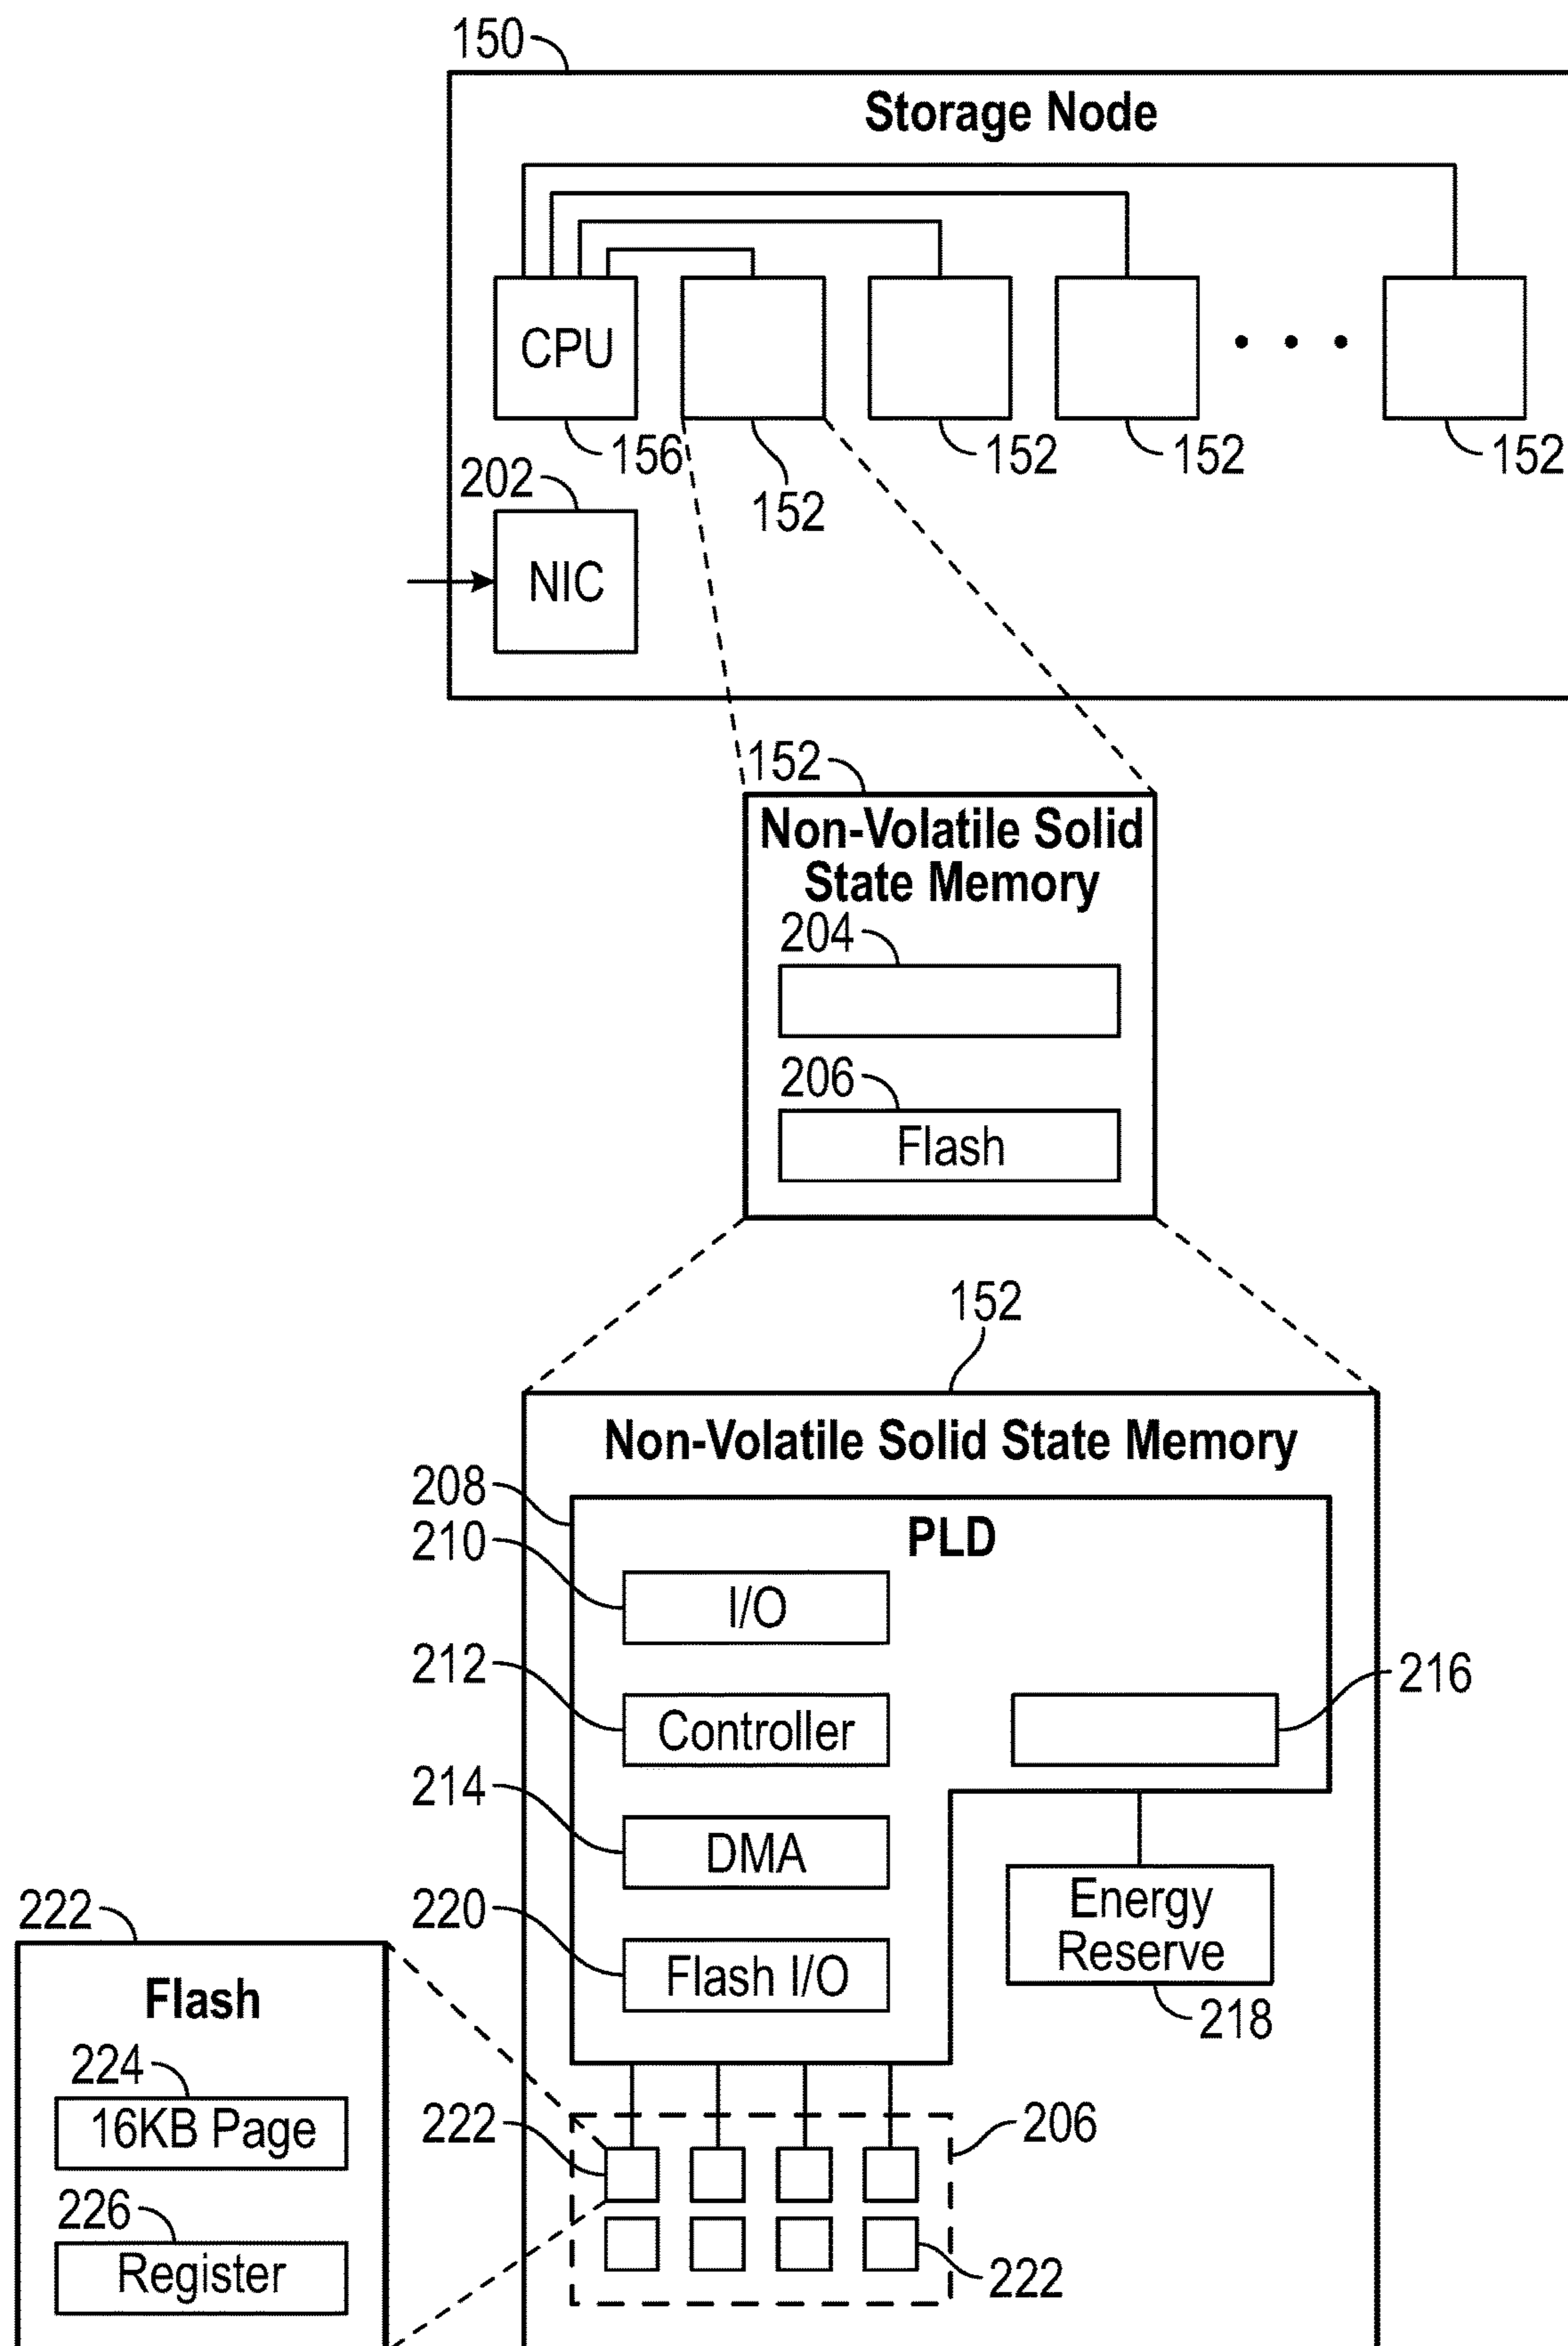


FIG. 3

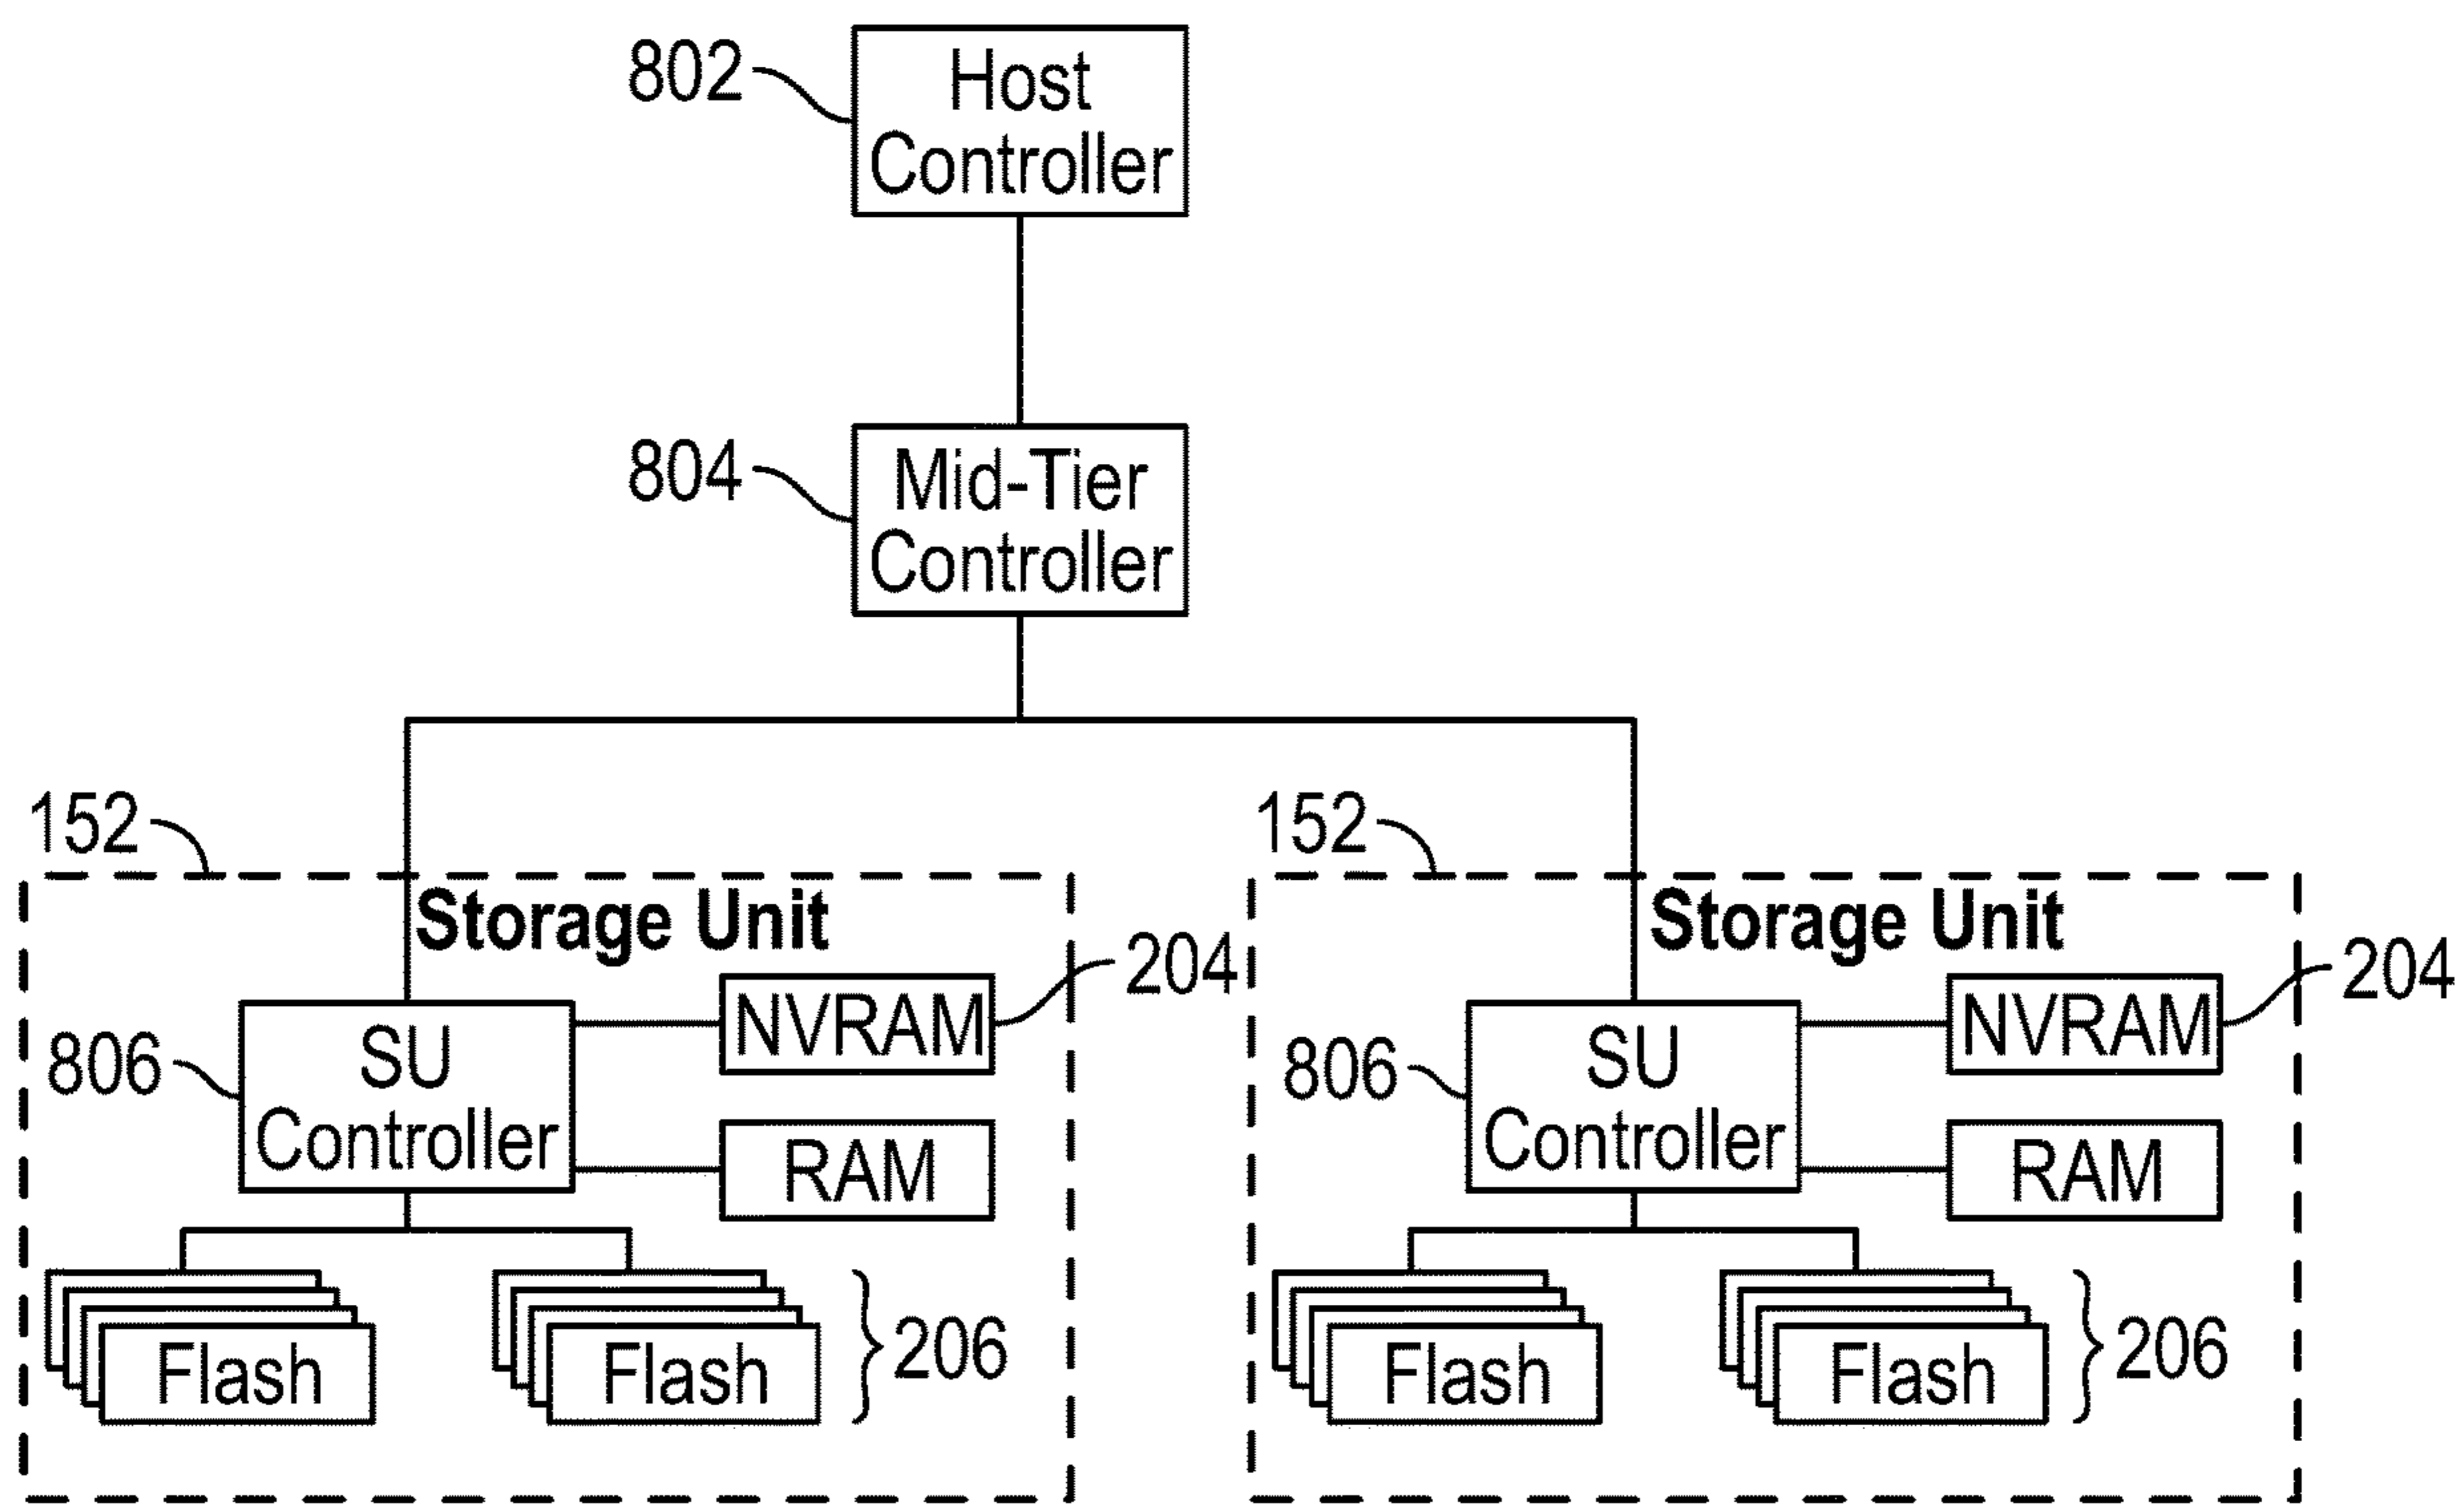


FIG. 4

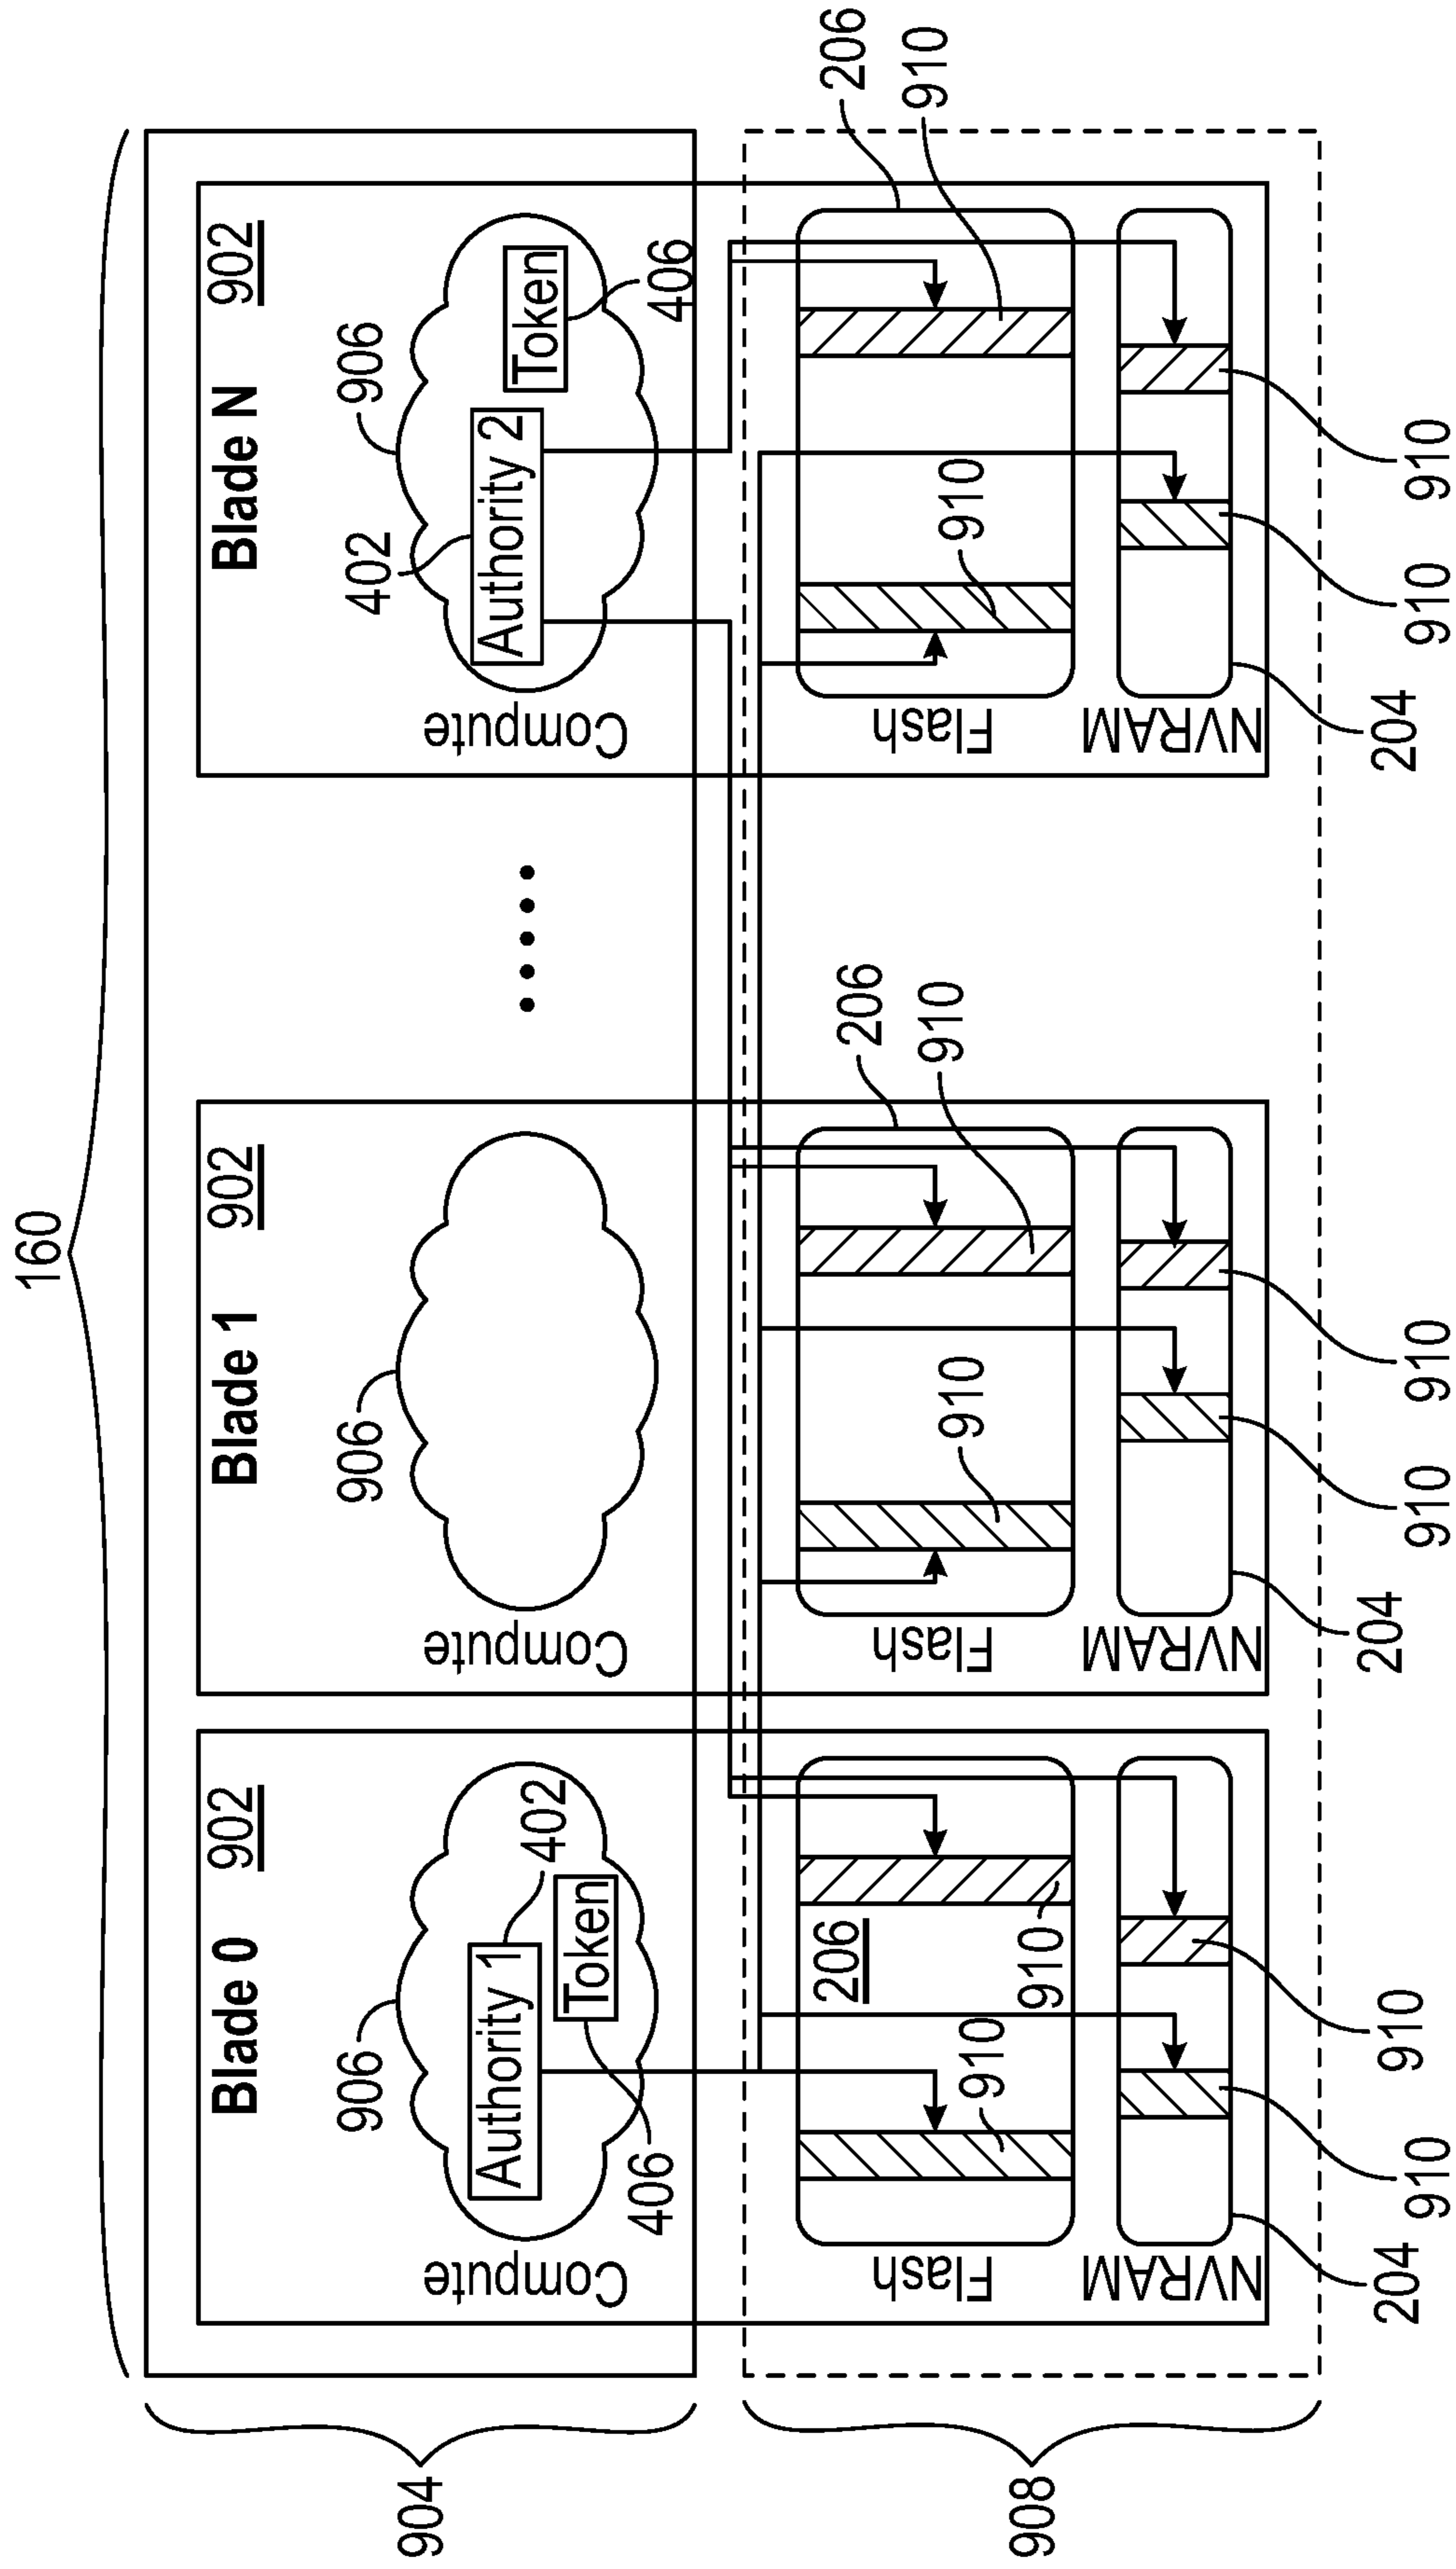


FIG. 5

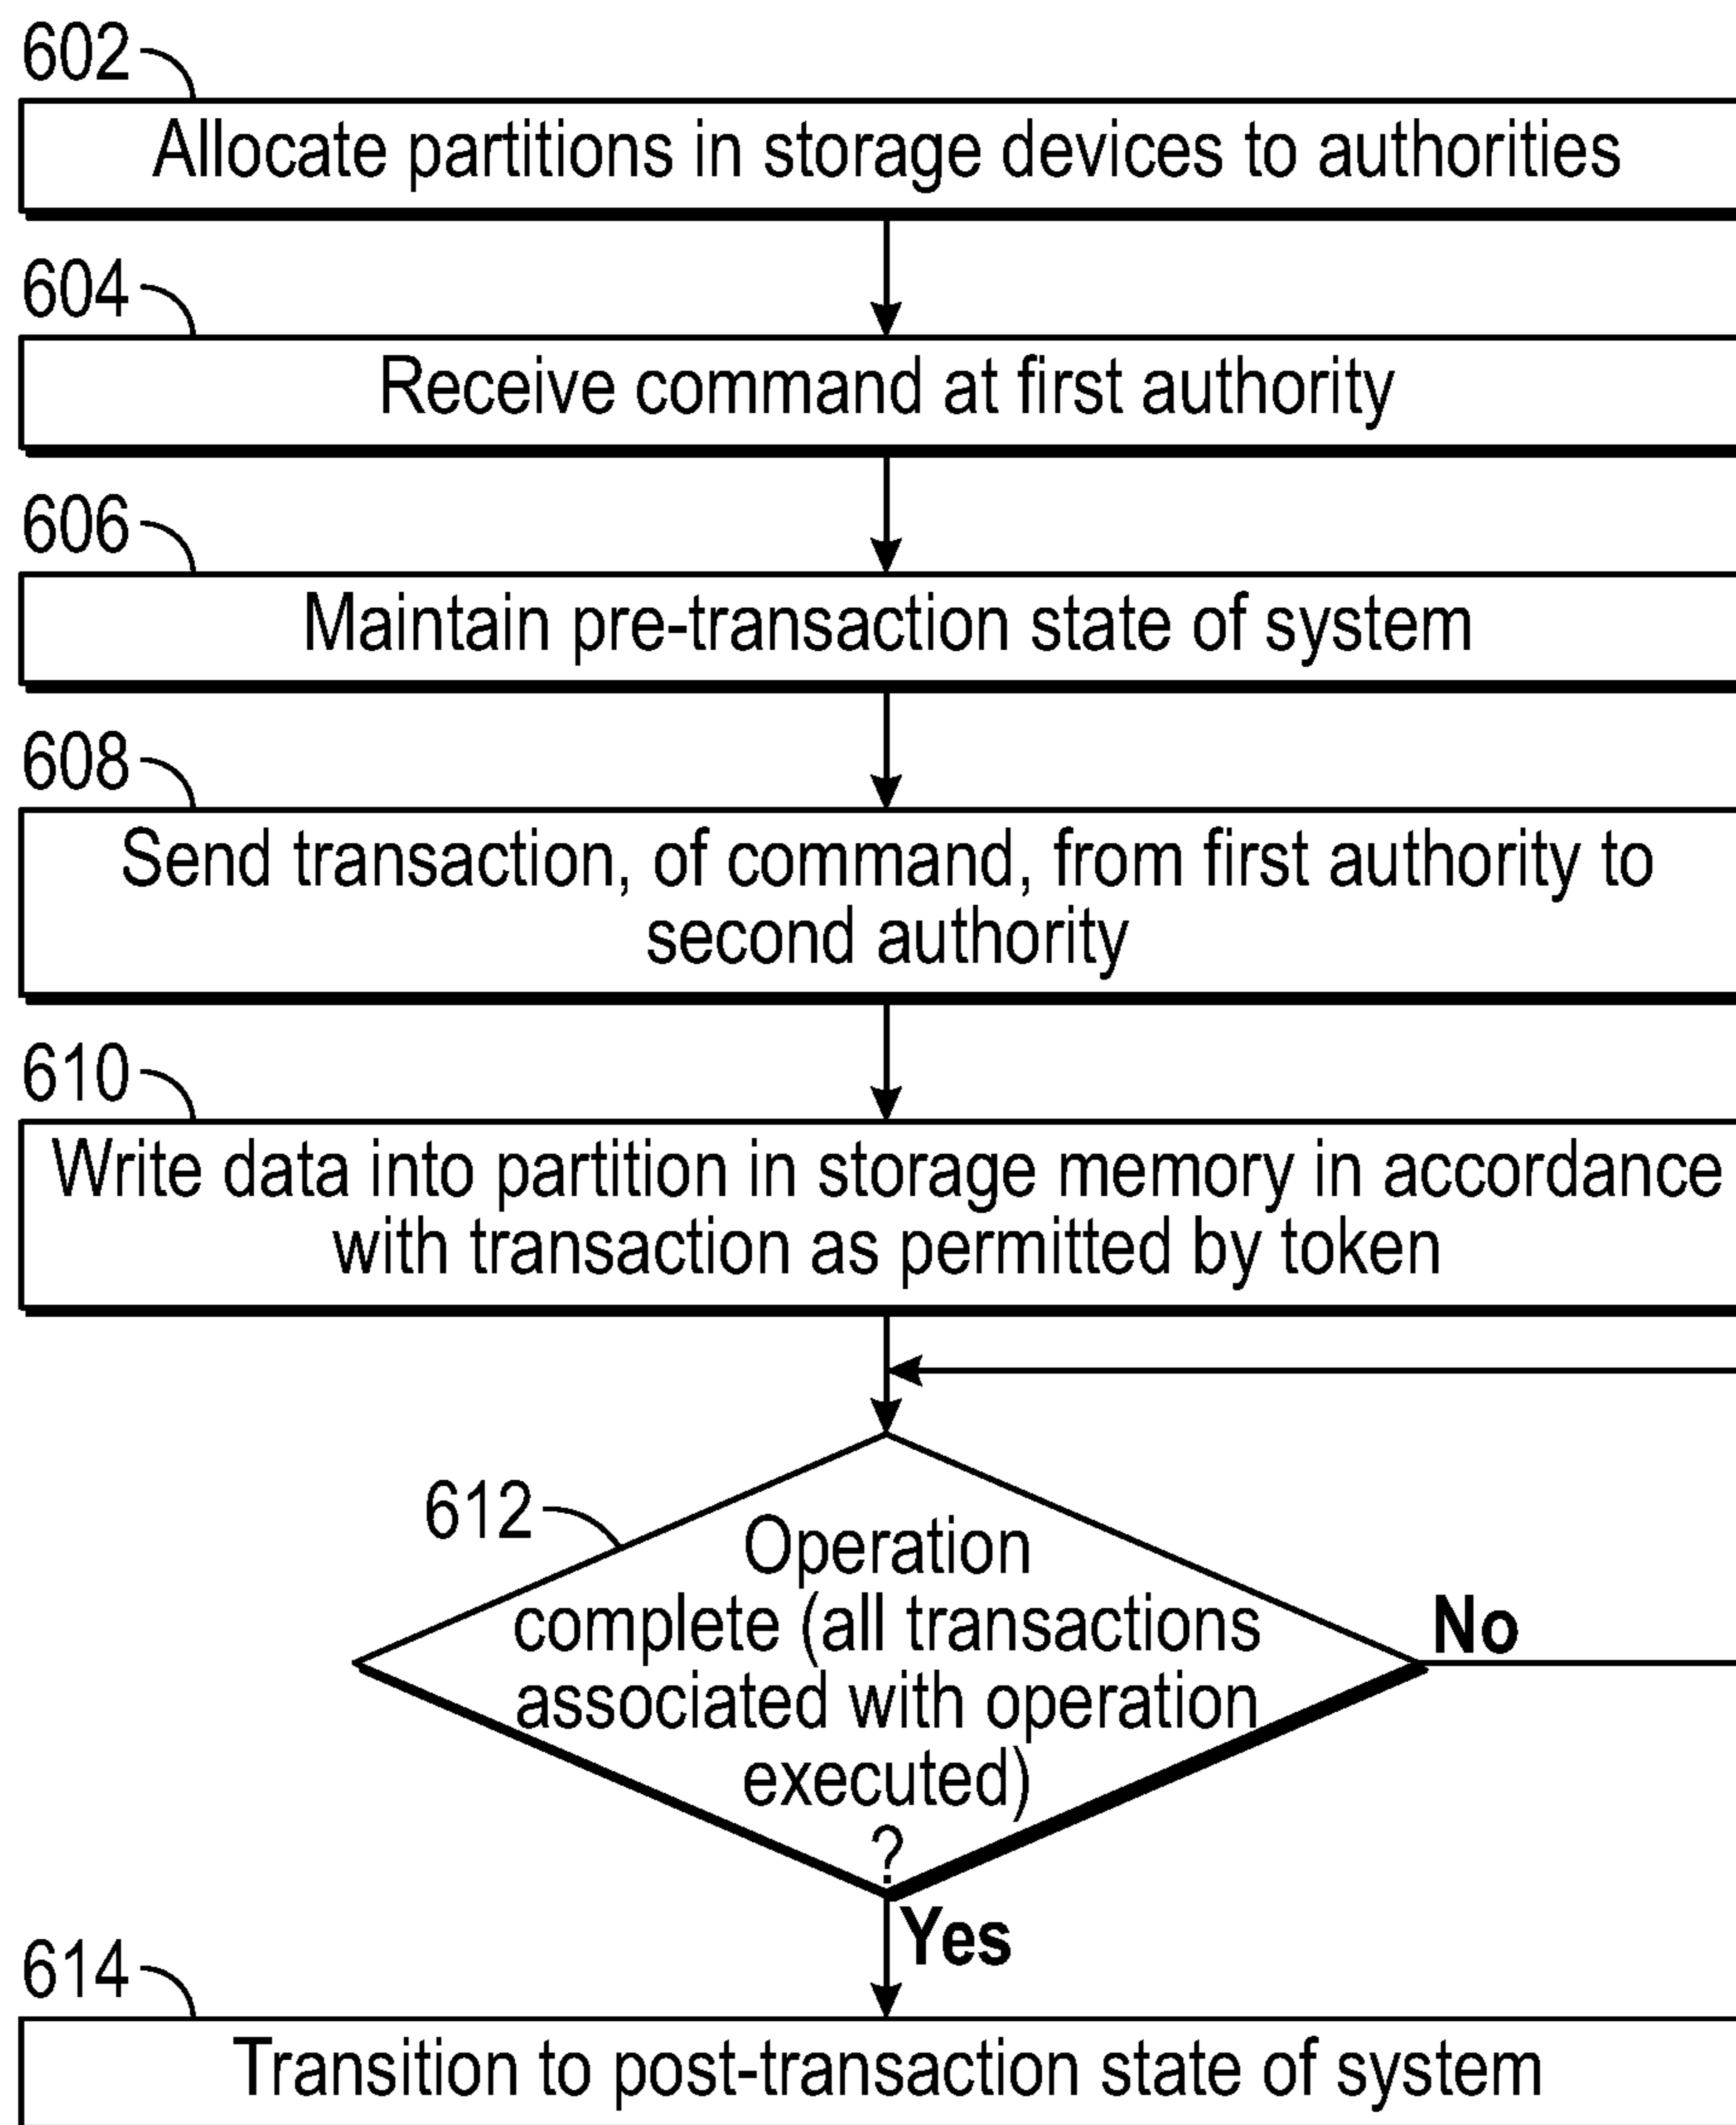


FIG. 6

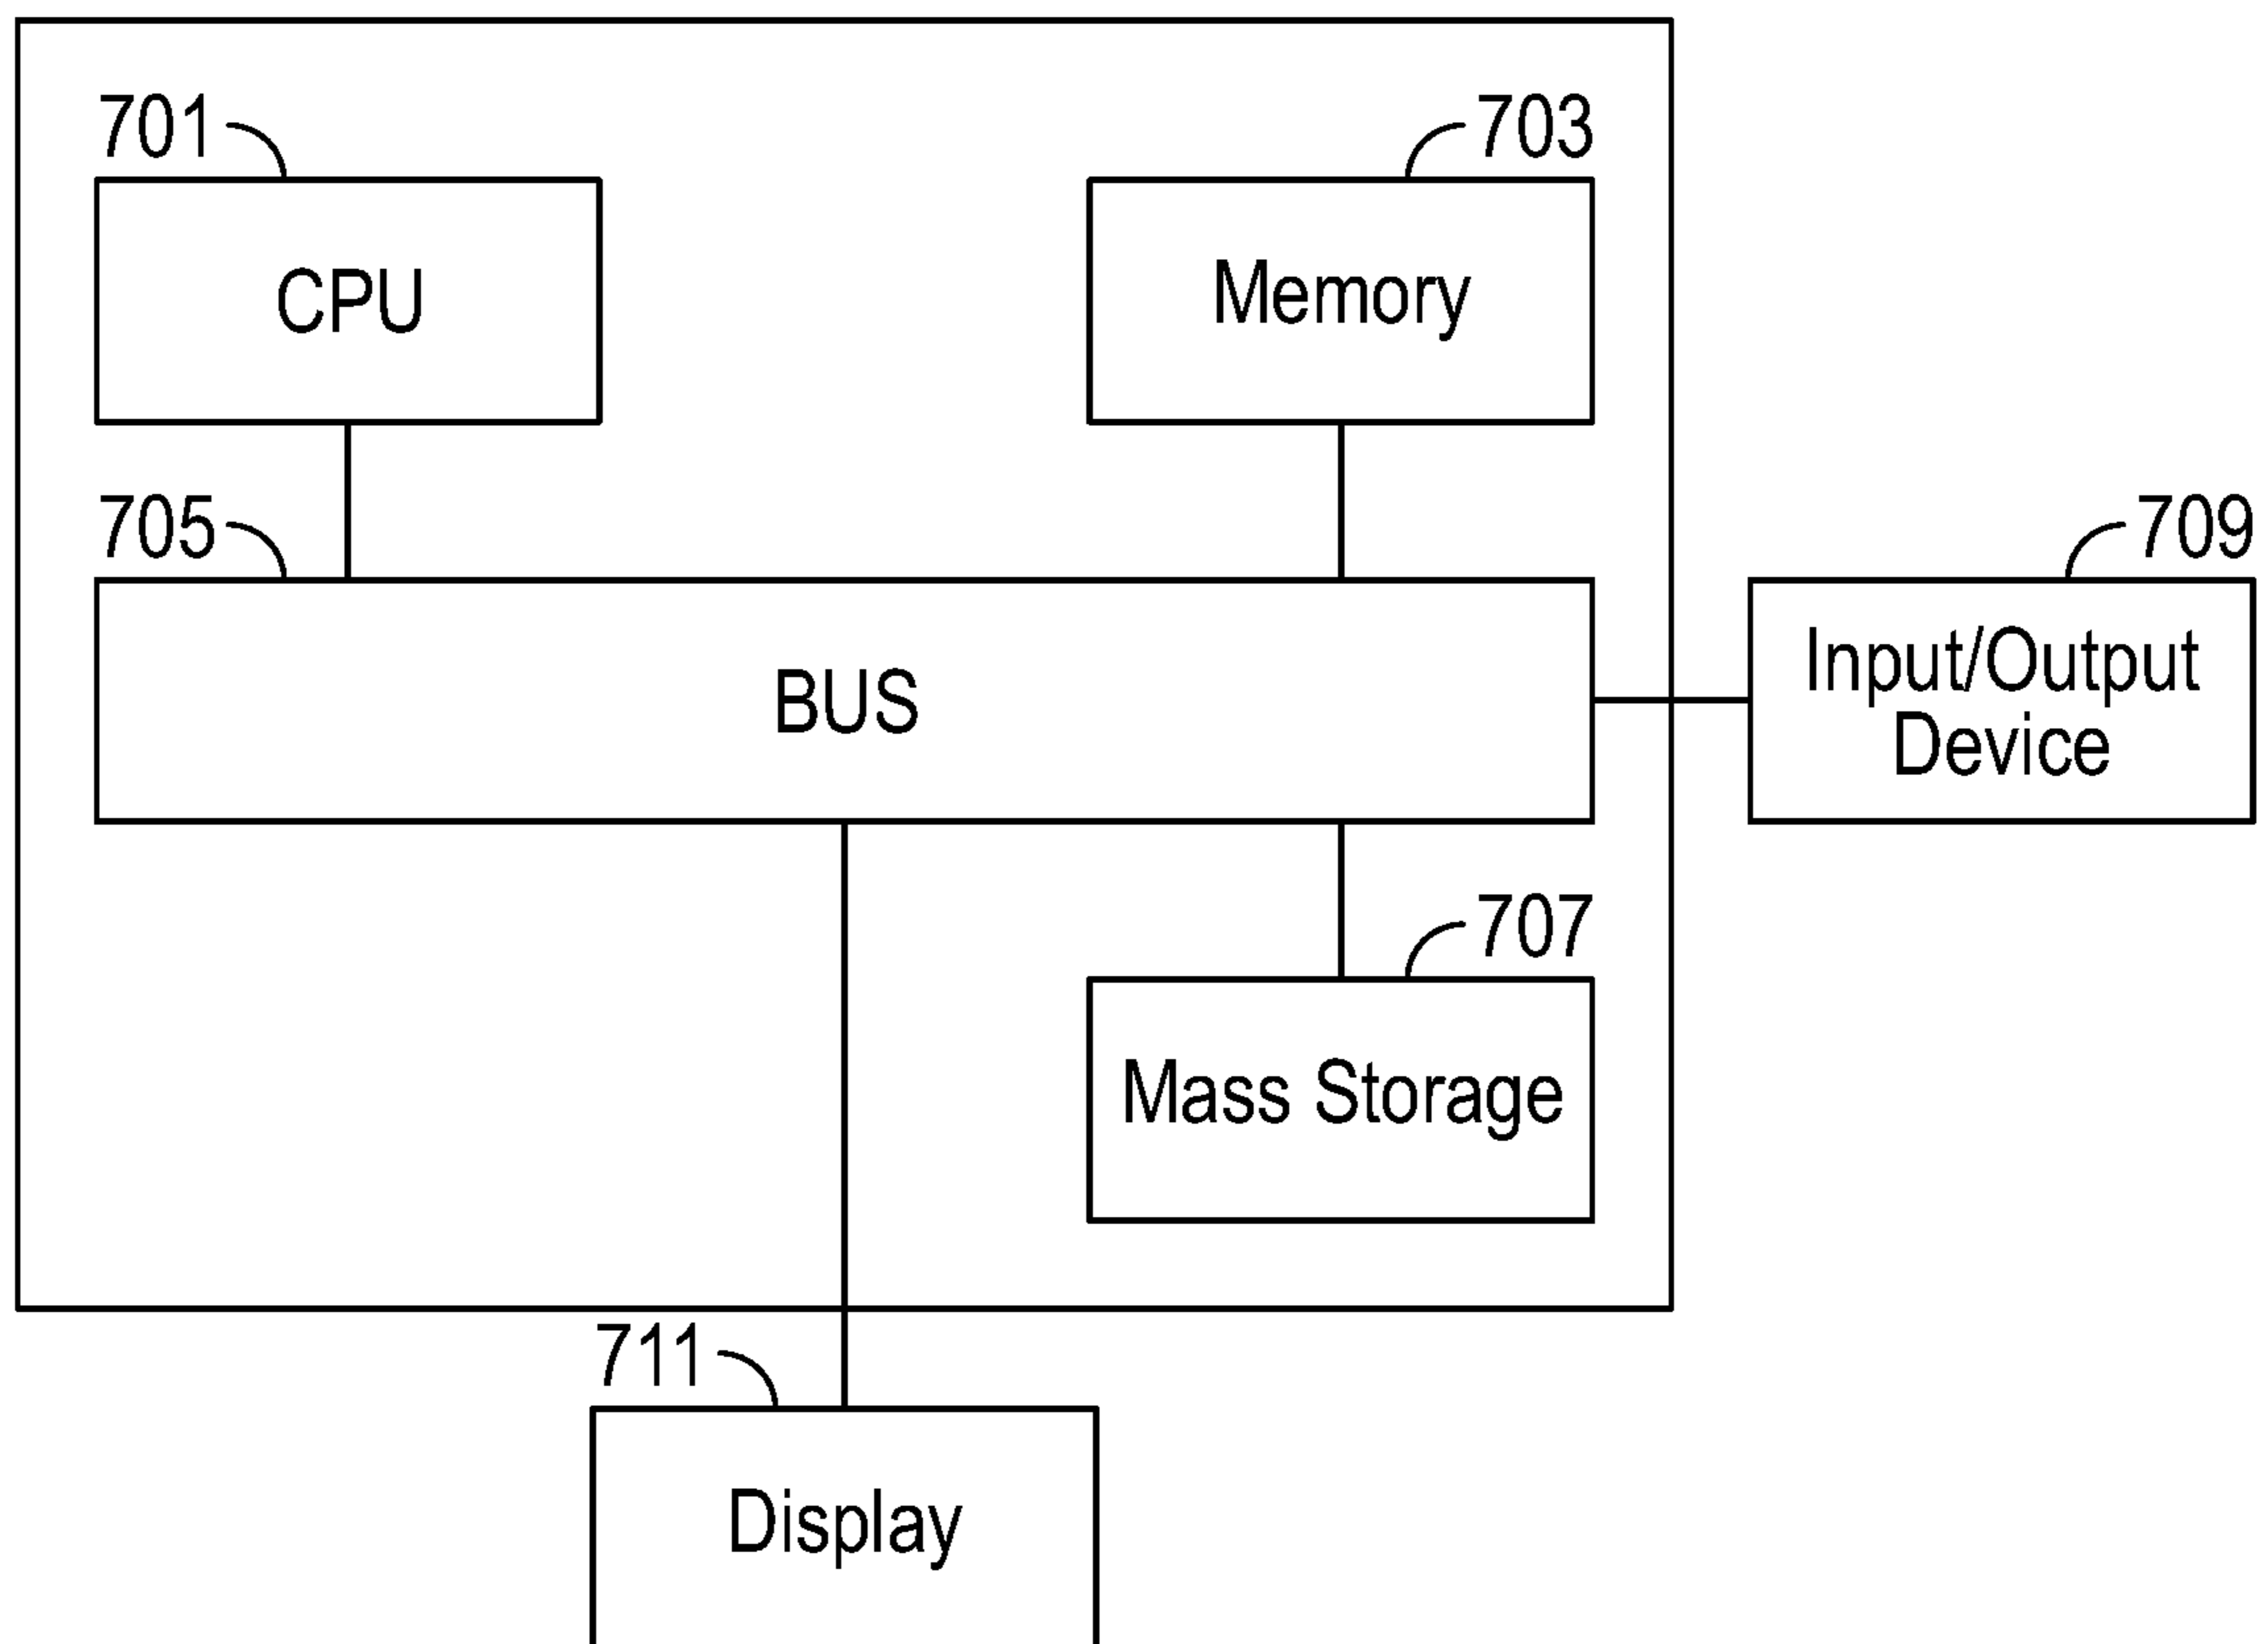


FIG. 7

1**TRANSACTION PROCESSING WITH
DIFFERING CAPACITY STORAGE**

BACKGROUND

Solid-state memory, such as flash, is currently in use in solid-state drives (SSD) to augment or replace conventional hard disk drives (HDD), writable CD (compact disk) or writable DVD (digital versatile disk) drives, collectively known as spinning media, and tape drives, for storage of large amounts of data. Flash and other solid-state memories have characteristics that differ from spinning media. Yet, many solid-state drives are designed to conform to hard disk drive standards for compatibility reasons, which makes it difficult to provide enhanced features or take advantage of unique aspects of flash and other solid-state memory. In a storage cluster environment, it may be difficult to maintain a consistent system state with distributed transaction and operation processing. One known mechanism for doing so, applying a global lock, during a transaction or operation, is effective but limits system bandwidth and throughput.

It is within this context that the embodiments arise.

SUMMARY

In some embodiments, a method of processing transactions associated with a command in a storage system is provided. The method includes receiving, at a first authority of the storage system, a command relating to user data. The method includes sending a transaction of the command, from the first authority to a second authority of the storage system, wherein a token accompanies or is associated with the transaction and writing data in accordance with the transaction as permitted by the token into a partition that is allocated to the second authority in a storage device of the storage system. A tangible, non-transitory, computer-readable media having instructions thereupon which, when executed by a processor, cause the processor to perform the method is also provided. A storage system comprising a memory and a processing unit having a processor configured to execute actions of the method is also provided.

Other aspects and advantages of the embodiments will become apparent from the following detailed description taken in conjunction with the accompanying drawings which illustrate, by way of example, the principles of the described embodiments.

BRIEF DESCRIPTION OF THE DRAWINGS

The described embodiments and the advantages thereof may best be understood by reference to the following description taken in conjunction with the accompanying drawings. These drawings in no way limit any changes in form and detail that may be made to the described embodiments by one skilled in the art without departing from the spirit and scope of the described embodiments.

FIG. 1 is a perspective view of a storage cluster with multiple storage nodes and internal storage coupled to each storage node to provide network attached storage, in accordance with some embodiments.

FIG. 2 is a block diagram showing an interconnect switch coupling multiple storage nodes in accordance with some embodiments.

FIG. 3 is a multiple level block diagram, showing contents of a storage node and contents of one of the non-volatile solid state storage units in accordance with some embodiments.

2

FIG. 4 shows a storage server environment, which uses embodiments of the storage nodes and storage units of FIGS. 1-3 in accordance with some embodiments.

FIG. 5 is a blade hardware block diagram, showing a control plane, compute and storage planes, and authorities interacting with underlying physical resources to perform distributed transactions with token-associated execution, in accordance with some embodiments.

FIG. 6 is a flow diagram of a method for processing transactions associated with an operation, which can be practiced in the storage cluster of FIGS. 1-5, in accordance with some embodiments.

FIG. 7 is an illustration showing an exemplary computing device which may implement the embodiments described herein.

DETAILED DESCRIPTION

The embodiments below describe a storage cluster that stores user data, such as user data originating from one or more user or client systems or other sources external to the storage cluster. The storage cluster distributes user data across storage nodes housed within a chassis, using erasure coding and redundant copies of metadata. Erasure coding refers to a method of data protection or reconstruction in which data is stored across a set of different locations, such as disks, storage nodes or geographic locations. Flash memory is one type of solid-state memory that may be integrated with the embodiments, although the embodiments may be extended to other types of solid-state memory or other storage medium, including non-solid state memory. Control of storage locations and workloads are distributed across the storage locations in a clustered peer-to-peer system. Tasks such as mediating communications between the various storage nodes, detecting when a storage node has become unavailable, and balancing I/Os (inputs and outputs) across the various storage nodes, are all handled on a distributed basis. Data is laid out or distributed across multiple storage nodes in data fragments or stripes that support data recovery in some embodiments. Ownership of data can be reassigned within a cluster, independent of input and output patterns. This architecture described in more detail below allows a storage node in the cluster to fail, with the system remaining operational, since the data can be reconstructed from other storage nodes and thus remain available for input and output operations. In various embodiments, a storage node may be referred to as a cluster node, a blade, or a server.

Authorities, implemented in storage nodes in the storage cluster, process transactions associated with operations. Each transaction is accompanied by or associated with a token, which gives permission to execute the transaction. The authorities are able to maintain a pre-transaction state of the system until completion of the operation. This is accomplished without a global lock across the system, and also enables restart of an operation in case of a disruption or other failure. Various system aspects are discussed below with reference to FIGS. 1-3. Execution of transactions, using tokens, is described with reference to FIGS. 4-6.

The storage cluster is contained within a chassis, i.e., an enclosure housing one or more storage nodes. A mechanism to provide power to each storage node, such as a power distribution bus, and a communication mechanism, such as a communication bus that enables communication between the storage nodes are included within the chassis. The storage cluster can run as an independent system in one location according to some embodiments. In one embodi-

ment, a chassis contains at least two instances of both the power distribution and the communication bus which may be enabled or disabled independently. The internal communication bus may be an Ethernet bus, however, other technologies such as Peripheral Component Interconnect (PCI) Express, InfiniBand, and others, are equally suitable. The chassis provides a port for an external communication bus for enabling communication between multiple chassis, directly or through a switch, and with client systems. The external communication may use a technology such as Ethernet, InfiniBand, Fibre Channel, etc. In some embodiments, the external communication bus uses different communication bus technologies for inter-chassis and client communication. If a switch is deployed within or between chassis, the switch may act as a translation between multiple protocols or technologies. When multiple chassis are connected to define a storage cluster, the storage cluster may be accessed by a client using either proprietary interfaces or standard interfaces such as network file system (NFS), common internet file system (CIFS), small computer system interface (SCSI) or hypertext transfer protocol (HTTP). Translation from the client protocol may occur at the switch, chassis external communication bus or within each storage node.

Each storage node may be one or more storage servers and each storage server is connected to one or more non-volatile solid state memory units, which may be referred to as storage units or storage devices. One embodiment includes a single storage server in each storage node and between one to eight non-volatile solid state memory units, however this one example is not meant to be limiting. The storage server may include a processor, dynamic random access memory (DRAM) and interfaces for the internal communication bus and power distribution for each of the power buses. Inside the storage node, the interfaces and storage unit share a communication bus, e.g., PCI Express, in some embodiments. The non-volatile solid state memory units may directly access the internal communication bus interface through a storage node communication bus, or request the storage node to access the bus interface. The non-volatile solid state memory unit contains an embedded central processing unit (CPU), solid state storage controller, and a quantity of solid state mass storage, e.g., between 2-32 terabytes (TB) in some embodiments. An embedded volatile storage medium, such as DRAM, and an energy reserve apparatus are included in the non-volatile solid state memory unit. In some embodiments, the energy reserve apparatus is a capacitor, super-capacitor, or battery that enables transferring a subset of DRAM contents to a stable storage medium in the case of power loss. In some embodiments, the non-volatile solid state memory unit is constructed with a storage class memory, such as phase change or magnetoresistive random access memory (MRAM) that substitutes for DRAM and enables a reduced power hold-up apparatus.

One of many features of the storage nodes and non-volatile solid state storage is the ability to proactively rebuild data in a storage cluster. The storage nodes and non-volatile solid state storage can determine when a storage node or non-volatile solid state storage in the storage cluster is unreachable, independent of whether there is an attempt to read data involving that storage node or non-volatile solid state storage. The storage nodes and non-volatile solid state storage then cooperate to recover and rebuild the data in at least partially new locations. This constitutes a proactive rebuild, in that the system rebuilds data without waiting until the data is needed for a read access initiated from a client

system employing the storage cluster. These and further details of the storage memory and operation thereof are discussed below.

FIG. 1 is a perspective view of a storage cluster 160, with multiple storage nodes 150 and internal solid-state memory coupled to each storage node to provide network attached storage or storage area network, in accordance with some embodiments. A network attached storage, storage area network, or a storage cluster, or other storage memory, could include one or more storage clusters 160, each having one or more storage nodes 150, in a flexible and reconfigurable arrangement of both the physical components and the amount of storage memory provided thereby. The storage cluster 160 is designed to fit in a rack, and one or more racks can be set up and populated as desired for the storage memory. The storage cluster 160 has a chassis 138 having multiple slots 142. It should be appreciated that chassis 138 may be referred to as a housing, enclosure, or rack unit. In one embodiment, the chassis 138 has fourteen slots 142, although other numbers of slots are readily devised. For example, some embodiments have four slots, eight slots, sixteen slots, thirty-two slots, or other suitable number of slots. Each slot 142 can accommodate one storage node 150 in some embodiments. Chassis 138 includes flaps 148 that can be utilized to mount the chassis 138 on a rack. Fans 144 provide air circulation for cooling of the storage nodes 150 and components thereof, although other cooling components could be used, or an embodiment could be devised without cooling components. A switch fabric 146 couples storage nodes 150 within chassis 138 together and to a network for communication to the memory. In an embodiment depicted in FIG. 1, the slots 142 to the left of the switch fabric 146 and fans 144 are shown occupied by storage nodes 150, while the slots 142 to the right of the switch fabric 146 and fans 144 are empty and available for insertion of storage node 150 for illustrative purposes. This configuration is one example, and one or more storage nodes 150 could occupy the slots 142 in various further arrangements. The storage node arrangements need not be sequential or adjacent in some embodiments. Storage nodes 150 are hot pluggable, meaning that a storage node 150 can be inserted into a slot 142 in the chassis 138, or removed from a slot 142, without stopping or powering down the system. Upon insertion or removal of storage node 150 from slot 142, the system automatically reconfigures in order to recognize and adapt to the change. Reconfiguration, in some embodiments, includes restoring redundancy and/or rebalancing data or load.

Each storage node 150 can have multiple components. In the embodiment shown here, the storage node 150 includes a printed circuit board 158 populated by a CPU 156, i.e., processor, a memory 154 coupled to the CPU 156, and a non-volatile solid state storage 152 coupled to the CPU 156, although other mountings and/or components could be used in further embodiments. The memory 154 has instructions which are executed by the CPU 156 and/or data operated on by the CPU 156. As further explained below, the non-volatile solid state storage 152 includes flash or, in further embodiments, other types of solid-state memory.

Referring to FIG. 1, storage cluster 160 is scalable, meaning that storage capacity with non-uniform storage sizes is readily added, as described above. One or more storage nodes 150 can be plugged into or removed from each chassis and the storage cluster self-configures in some embodiments. Plug-in storage nodes 150, whether installed in a chassis as delivered or later added, can have different sizes. For example, in one embodiment a storage node 150

5

can have any multiple of 4 TB, e.g., 8 TB, 12 TB, 16 TB, 32 TB, etc. In further embodiments, a storage node **150** could have any multiple of other storage amounts or capacities. Storage capacity of each storage node **150** is broadcast, and influences decisions of how to stripe the data. For maximum storage efficiency, an embodiment can self-configure as wide as possible in the stripe, subject to a predetermined requirement of continued operation with loss of up to one, or up to two, non-volatile solid state storage units **152** or storage nodes **150** within the chassis.

FIG. 2 is a block diagram showing a communications interconnect **170** and power distribution bus **172** coupling multiple storage nodes **150**. Referring back to FIG. 1, the communications interconnect **170** can be included in or implemented with the switch fabric **146** in some embodiments. Where multiple storage clusters **160** occupy a rack, the communications interconnect **170** can be included in or implemented with a top of rack switch, in some embodiments. As illustrated in FIG. 2, storage cluster **160** is enclosed within a single chassis **138**. External port **176** is coupled to storage nodes **150** through communications interconnect **170**, while external port **174** is coupled directly to a storage node. External power port **178** is coupled to power distribution bus **172**. Storage nodes **150** may include varying amounts and differing capacities of non-volatile solid state storage **152** as described with reference to FIG. 1. In addition, one or more storage nodes **150** may be a compute only storage node as illustrated in FIG. 2. Authorities **168** are implemented on the non-volatile solid state storages **152**, for example as lists or other data structures stored in memory. In some embodiments the authorities are stored within the non-volatile solid state storage **152** and supported by software executing on a controller or other processor of the non-volatile solid state storage **152**. In a further embodiment, authorities **168** are implemented on the storage nodes **150**, for example as lists or other data structures stored in the memory **154** and supported by software executing on the CPU **156** of the storage node **150**. Authorities **168** control how and where data is stored in the non-volatile solid state storages **152** in some embodiments. This control assists in determining which type of erasure coding scheme is applied to the data, and which storage nodes **150** have which portions of the data. Each authority **168** may be assigned to a non-volatile solid state storage **152**. Each authority may control a range of inode numbers, segment numbers, or other data identifiers which are assigned to data by a file system, by the storage nodes **150**, or by the non-volatile solid state storage **152**, in various embodiments.

Every piece of data, and every piece of metadata, has redundancy in the system in some embodiments. In addition, every piece of data and every piece of metadata has an owner, which may be referred to as an authority. If that authority is unreachable, for example through failure of a storage node, there is a plan of succession for how to find that data or that metadata. In various embodiments, there are redundant copies of authorities **168**. Authorities **168** have a relationship to storage nodes **150** and non-volatile solid state storage **152** in some embodiments. Each authority **168**, covering a range of data segment numbers or other identifiers of the data, may be assigned to a specific non-volatile solid state storage **152**. In some embodiments the authorities **168** for all of such ranges are distributed over the non-volatile solid state storages **152** of a storage cluster. Each storage node **150** has a network port that provides access to the non-volatile solid state storage(s) **152** of that storage node **150**. Data can be stored in a segment, which is associated with a segment number and that segment number

6

is an indirection for a configuration of a RAID (redundant array of independent disks) stripe in some embodiments. The assignment and use of the authorities **168** thus establishes an indirection to data. Indirection may be referred to as the ability to reference data indirectly, in this case via an authority **168**, in accordance with some embodiments. A segment identifies a set of non-volatile solid state storage **152** and a local identifier into the set of non-volatile solid state storage **152** that may contain data. In some embodiments, the local identifier is an offset into the device and may be reused sequentially by multiple segments. In other embodiments the local identifier is unique for a specific segment and never reused. The offsets in the non-volatile solid state storage **152** are applied to locating data for writing to or reading from the non-volatile solid state storage **152** (in the form of a RAID stripe). Data is striped across multiple units of non-volatile solid state storage **152**, which may include or be different from the non-volatile solid state storage **152** having the authority **168** for a particular data segment.

If there is a change in where a particular segment of data is located, e.g., during a data move or a data reconstruction, the authority **168** for that data segment should be consulted, at that non-volatile solid state storage **152** or storage node **150** having that authority **168**. In order to locate a particular piece of data, embodiments calculate a hash value for a data segment or apply an inode number or a data segment number. The output of this operation points to a non-volatile solid state storage **152** having the authority **168** for that particular piece of data. In some embodiments there are two stages to this operation. The first stage maps an entity identifier (ID), e.g., a segment number, inode number, or directory number to an authority identifier. This mapping may include a calculation such as a hash or a bit mask. The second stage is mapping the authority identifier to a particular non-volatile solid state storage **152**, which may be done through an explicit mapping. The operation is repeatable, so that when the calculation is performed, the result of the calculation repeatably and reliably points to a particular non-volatile solid state storage **152** having that authority **168**. The operation may include the set of reachable storage nodes as input. If the set of reachable non-volatile solid state storage units changes the optimal set changes. In some embodiments, the persisted value is the current assignment (which is always true) and the calculated value is the target assignment the cluster will attempt to reconfigure towards. This calculation may be used to determine the optimal non-volatile solid state storage **152** for an authority in the presence of a set of non-volatile solid state storage **152** that are reachable and constitute the same cluster. The calculation also determines an ordered set of peer non-volatile solid state storage **152** that will also record the authority to non-volatile solid state storage mapping so that the authority may be determined even if the assigned non-volatile solid state storage is unreachable. A duplicate or substitute authority **168** may be consulted if a specific authority **168** is unavailable in some embodiments.

With reference to FIGS. 1 and 2, two of the many tasks of the CPU **156** on a storage node **150** are to break up write data, and reassemble read data. When the system has determined that data is to be written, the authority **168** for that data is located as above. When the segment ID for data is already determined the request to write is forwarded to the non-volatile solid state storage **152** currently determined to be the host of the authority **168** determined from the segment. The host CPU **156** of the storage node **150**, on which the non-volatile solid state storage **152** and corre-

sponding authority **168** reside, then breaks up or shards the data and transmits the data out to various non-volatile solid state storage **152**. The transmitted data is written as a data stripe in accordance with an erasure coding scheme. In some embodiments, data is requested to be pulled, and in other 5 embodiments, data is pushed. In reverse, when data is read, the authority **168** for the segment ID containing the data is located as described above. The host CPU **156** of the storage node **150** on which the non-volatile solid state storage **152** and corresponding authority **168** reside requests the data 10 from the non-volatile solid state storage and corresponding storage nodes pointed to by the authority. In some embodiments the data is read from flash storage as a data stripe. The host CPU **156** of storage node **150** then reassembles the read data, correcting any errors (if present) according to the 15 appropriate erasure coding scheme, and forwards the reassembled data to the network. In further embodiments, some or all of these tasks can be handled in the non-volatile solid state storage **152**. In some embodiments, the segment host requests the data be sent to storage node **150** by requesting 20 pages from storage and then sending the data to the storage node making the original request.

In some systems, for example in UNIX-style file systems, data is handled with an index node or inode, which specifies a data structure that represents an object in a file system. The 25 object could be a file or a directory, for example. Metadata may accompany the object, as attributes such as permission data and a creation timestamp, among other attributes. A segment number could be assigned to all or a portion of such an object in a file system. In other systems, data segments 30 are handled with a segment number assigned elsewhere. For purposes of discussion, the unit of distribution is an entity, and an entity can be a file, a directory or a segment. That is, entities are units of data or metadata stored by a storage system. Entities are grouped into sets called authorities. 35 Each authority has an authority owner, which is a storage node that has the exclusive right to update the entities in the authority. In other words, a storage node contains the authority, and that the authority, in turn, contains entities.

A segment is a logical container of data in accordance 40 with some embodiments. A segment is an address space between medium address space and physical flash locations, i.e., the data segment number, are in this address space. Segments may also contain meta-data, which enable data redundancy to be restored (rewritten to different flash loca- 45 tions or devices) without the involvement of higher level software. In one embodiment, an internal format of a segment contains client data and medium mappings to determine the position of that data. Each data segment is protected, e.g., from memory and other failures, by breaking the 50 segment into a number of data and parity shards, where applicable. The data and parity shards are distributed, i.e., striped, across non-volatile solid state storage **152** coupled to the host CPUs **156** (See FIG. 5) in accordance with an erasure coding scheme. Usage of the term segments refers to 55 the container and its place in the address space of segments in some embodiments. Usage of the term stripe refers to the same set of shards as a segment and includes how the shards are distributed along with redundancy or parity information in accordance with some embodiments.

A series of address-space transformations takes place across an entire storage system. At the top are the directory entries (file names) which link to an inode. Modes point into 60 medium address space, where data is logically stored. Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots.

Medium addresses may be mapped through a series of indirect mediums to spread the load of large files, or implement data services like deduplication or snapshots. Segment addresses are then translated into physical flash 5 locations. Physical flash locations have an address range bounded by the amount of flash in the system in accordance with some embodiments. Medium addresses and segment addresses are logical containers, and in some embodiments use a 128 bit or larger identifier so as to be practically 10 infinite, with a likelihood of reuse calculated as longer than the expected life of the system. Addresses from logical containers are allocated in a hierarchical fashion in some embodiments. Initially, each non-volatile solid state storage unit **152** may be assigned a range of address space. Within 15 this assigned range, the non-volatile solid state storage **152** is able to allocate addresses without synchronization with other non-volatile solid state storage **152**.

Data and metadata is stored by a set of underlying storage layouts that are optimized for varying workload patterns and 20 storage devices. These layouts incorporate multiple redundancy schemes, compression formats and index algorithms. Some of these layouts store information about authorities and authority masters, while others store file metadata and file data. The redundancy schemes include error correction 25 codes that tolerate corrupted bits within a single storage device (such as a NAND flash chip), erasure codes that tolerate the failure of multiple storage nodes, and replication schemes that tolerate data center or regional failures. In some embodiments, low density parity check (LDPC) code 30 is used within a single storage unit. Reed-Solomon encoding is used within a storage cluster, and mirroring is used within a storage grid in some embodiments. Metadata may be stored using an ordered log structured index (such as a Log Structured Merge Tree), and large data may not be stored in 35 a log structured layout.

In order to maintain consistency across multiple copies of an entity, the storage nodes agree implicitly on two things through calculations: (1) the authority that contains the 40 entity, and (2) the storage node that contains the authority. The assignment of entities to authorities can be done by pseudo randomly assigning entities to authorities, by splitting entities into ranges based upon an externally produced key, or by placing a single entity into each authority. Examples of pseudorandom schemes are linear hashing and 45 the Replication Under Scalable Hashing (RUSH) family of hashes, including Controlled Replication Under Scalable Hashing (CRUSH). In some embodiments, pseudo-random assignment is utilized only for assigning authorities to nodes because the set of nodes can change. The set of authorities 50 cannot change so any subjective function may be applied in these embodiments. Some placement schemes automatically place authorities on storage nodes, while other placement schemes rely on an explicit mapping of authorities to storage nodes. In some embodiments, a pseudorandom scheme is 55 utilized to map from each authority to a set of candidate authority owners. A pseudorandom data distribution function related to CRUSH may assign authorities to storage nodes and create a list of where the authorities are assigned. Each storage node has a copy of the pseudorandom data 60 distribution function, and can arrive at the same calculation for distributing, and later finding or locating an authority. Each of the pseudorandom schemes requires the reachable set of storage nodes as input in some embodiments in order to conclude the same target nodes. Once an entity has been 65 placed in an authority, the entity may be stored on physical devices so that no expected failure will lead to unexpected data loss. In some embodiments, rebalancing algorithms

attempt to store the copies of all entities within an authority in the same layout and on the same set of machines.

Examples of expected failures include device failures, stolen machines, datacenter fires, and regional disasters, such as nuclear or geological events. Different failures lead to different levels of acceptable data loss. In some embodiments, a stolen storage node impacts neither the security nor the reliability of the system, while depending on system configuration, a regional event could lead to no loss of data, a few seconds or minutes of lost updates, or even complete data loss.

In the embodiments, the placement of data for storage redundancy is independent of the placement of authorities for data consistency. In some embodiments, storage nodes that contain authorities do not contain any persistent storage. Instead, the storage nodes are connected to non-volatile solid state storage units that do not contain authorities. The communications interconnect between storage nodes and non-volatile solid state storage units consists of multiple communication technologies and has non-uniform performance and fault tolerance characteristics. In some embodiments, as mentioned above, non-volatile solid state storage units are connected to storage nodes via PCI express, storage nodes are connected together within a single chassis using Ethernet backplane, and chassis are connected together to form a storage cluster. Storage clusters are connected to clients using Ethernet or fiber channel in some embodiments. If multiple storage clusters are configured into a storage grid, the multiple storage clusters are connected using the Internet or other long-distance networking links, such as a "metro scale" link or private link that does not traverse the internet.

Authority owners have the exclusive right to modify entities, to migrate entities from one non-volatile solid state storage unit to another non-volatile solid state storage unit, and to add and remove copies of entities. This allows for maintaining the redundancy of the underlying data. When an authority owner fails, is going to be decommissioned, or is overloaded, the authority is transferred to a new storage node. Transient failures make it non-trivial to ensure that all non-faulty machines agree upon the new authority location. The ambiguity that arises due to transient failures can be achieved automatically by a consensus protocol such as Paxos, hot-warm failover schemes, via manual intervention by a remote system administrator, or by a local hardware administrator (such as by physically removing the failed machine from the cluster, or pressing a button on the failed machine). In some embodiments, a consensus protocol is used, and failover is automatic. If too many failures or replication events occur in too short a time period, the system goes into a self-preservation mode and halts replication and data movement activities until an administrator intervenes in accordance with some embodiments.

In some embodiments, the virtualized addresses are stored with sufficient redundancy. A continuous monitoring system correlates hardware and software status and the hardware identifiers. This allows detection and prediction of failures due to faulty components and manufacturing details. The monitoring system also enables the proactive transfer of authorities and entities away from impacted devices before failure occurs by removing the component from the critical path in some embodiments.

FIG. 3 is a multiple level block diagram, showing contents of a storage node **150** and contents of a non-volatile solid state storage **152** of the storage node **150**. Data is communicated to and from the storage node **150** by a network interface controller (NIC) **202** in some embodi-

ments. Each storage node **150** has a CPU **156**, and one or more non-volatile solid state storage **152**, as discussed above. Moving down one level in FIG. 3, each non-volatile solid state storage **152** has a relatively fast non-volatile solid state memory, such as nonvolatile random access memory (NVRAM) **204**, and flash memory **206**. In some embodiments, NVRAM **204** may be a component that does not require program/erase cycles (DRAM, MRAM, PCM), and can be a memory that can support being written vastly more often than the memory is read from. Moving down another level in FIG. 3, the NVRAM **204** is implemented in one embodiment as high speed volatile memory, such as dynamic random access memory (DRAM) **216**, backed up by energy reserve **218**. Energy reserve **218** provides sufficient electrical power to keep the DRAM **216** powered long enough for contents to be transferred to the flash memory **206** in the event of power failure. In some embodiments, energy reserve **218** is a capacitor, super-capacitor, battery, or other device, that supplies a suitable supply of energy sufficient to enable the transfer of the contents of DRAM **216** to a stable storage medium in the case of power loss. The flash memory **206** is implemented as multiple flash dies **222**, which may be referred to as packages of flash dies **222** or an array of flash dies **222**. It should be appreciated that the flash dies **222** could be packaged in any number of ways, with a single die per package, multiple dies per package (i.e. multichip packages), in hybrid packages, as bare dies on a printed circuit board or other substrate, as encapsulated dies, etc. In the embodiment shown, the non-volatile solid state storage **152** has a controller **212** or other processor, and an input output (I/O) port **210** coupled to the controller **212**. I/O port **210** is coupled to the CPU **156** and/or the network interface controller **202** of the flash storage node **150**. Flash input output (I/O) port **220** is coupled to the flash dies **222**, and a direct memory access unit (DMA) **214** is coupled to the controller **212**, the DRAM **216** and the flash dies **222**. In the embodiment shown, the I/O port **210**, controller **212**, DMA unit **214** and flash I/O port **220** are implemented on a programmable logic device (PLD) **208**, e.g., a field programmable gate array (FPGA). In this embodiment, each flash die **222** has pages, organized as sixteen kB (kilobyte) pages **224**, and a register **226** through which data can be written to or read from the flash die **222**. In further embodiments, other types of solid-state memory are used in place of, or in addition to flash memory illustrated within flash die **222**.

Storage clusters **160**, in various embodiments as disclosed herein, can be contrasted with storage arrays in general. The storage nodes **150** are part of a collection that creates the storage cluster **160**. Each storage node **150** owns a slice of data and computing required to provide the data. Multiple storage nodes **150** cooperate to store and retrieve the data. Storage memory or storage devices, as used in storage arrays in general, are less involved with processing and manipulating the data. Storage memory or storage devices in a storage array receive commands to read, write, or erase data. The storage memory or storage devices in a storage array are not aware of a larger system in which they are embedded, or what the data means. Storage memory or storage devices in storage arrays can include various types of storage memory, such as RAM, solid state drives, hard disk drives, etc. The storage units **152** described herein have multiple interfaces active simultaneously and serving multiple purposes. In some embodiments, some of the functionality of a storage node **150** is shifted into a storage unit **152**, transforming the storage unit **152** into a combination of storage unit **152** and storage node **150**. Placing computing (relative to storage

data) into the storage unit **152** places this computing closer to the data itself. The various system embodiments have a hierarchy of storage node layers with different capabilities. By contrast, in a storage array, a controller owns and knows everything about all of the data that the controller manages in a shelf or storage devices. In a storage cluster **160**, as described herein, multiple controllers in multiple storage units **152** and/or storage nodes **150** cooperate in various ways (e.g., for erasure coding, data sharding, metadata communication and redundancy, storage capacity expansion or contraction, data recovery, and so on).

FIG. **4** shows a storage server environment, which uses embodiments of the storage nodes **150** and storage units **152** of FIGS. **1-3**. In this version, each storage unit **152** has a processor such as controller **212** (see FIG. **3**), an FPGA (field programmable gate array), flash memory **206**, and NVRAM **204** (which is super-capacitor backed DRAM **216**, see FIGS. **2** and **3**) on a PCIe (peripheral component interconnect express) board in a chassis **138** (see FIG. **1**). The storage unit **152** may be implemented as a single board containing storage, and may be the largest tolerable failure domain inside the chassis. In some embodiments, up to two storage units **152** may fail and the device will continue with no data loss.

The physical storage is divided into named regions based on application usage in some embodiments. The NVRAM **204** is a contiguous block of reserved memory in the storage unit **152** DRAM **216**, and is backed by NAND flash. NVRAM **204** is logically divided into multiple memory regions written for two as spool (e.g., spool_region). Space within the NVRAM **204** spools is managed by each authority **402** independently. Each device provides an amount of storage space to each authority **402**. That authority **402** further manages lifetimes and allocations within that space. Examples of a spool include distributed transactions or notions. When the primary power to a storage unit **152** fails, onboard super-capacitors provide a short duration of power hold up. During this holdup interval, the contents of the NVRAM **204** are flushed to flash memory **206**. On the next power-on, the contents of the NVRAM **204** are recovered from the flash memory **206**.

As for the storage unit controller, the responsibility of the logical “controller” is distributed across each of the blades containing authorities **402**. This distribution of logical control is shown in FIG. **4** as a host controller **802**, mid-tier controller **804** and storage unit controller(s) **806**. Management of the control plane and the storage plane are treated independently, although parts may be physically co-located on the same blade. Each authority **402** effectively serves as an independent controller. Each authority **402** provides its own data and metadata structures, its own background workers, and maintains its own lifecycle.

FIG. **5** is a blade **902** hardware block diagram, showing a control plane **904**, compute and storage planes **906**, **908**, and authorities **402** interacting with underlying physical resources to perform distributed transactions with token-associated execution, using embodiments of the storage nodes **150** and storage units **152** of FIGS. **1-3** in the storage server environment of FIG. **4**. The control plane **904** is partitioned into a number of authorities **402** which can use the compute resources in the compute plane **906** to run on any of the blades **902**. The storage plane **908** is partitioned into a set of devices, each of which provides access to flash **206** and NVRAM **204** resources.

In the compute and storage planes **906**, **908** of FIG. **5**, the authorities **402** interact with the underlying physical resources (i.e., devices). From the point of view of an

authority **402**, its resources are striped over all of the physical devices. From the point of view of a device, it provides resources to all authorities **402**, irrespective of where the authorities happen to run. In order to communicate and represent the ownership of an authority **402**, including the right to record persistent changes on behalf of that authority **402**, the authority **402** must provide some evidence of authority ownership that can be independently verifiable. A token **406** is employed for this purpose and function in one embodiment.

One possibility for representing the authority token **406** would be to globally propagate each renewed token **406** out to each server and device in the storage cluster **160**. Doing so would require a prohibitively large number of communications that would scale proportionally with the number of servers. Additionally, pushing active updates to globally communicate new tokens **406** presents challenges for ensuring that the authority ownership can be safely advanced in the absence of positive communication (e.g., in the case of partial network partitioning). A solution to the above problems and constraints is to establish that an authority ownership token **406** is a time-bounded lease. The authority ownership and duration of the lease are represented by expressing the expiration time of the token **406**. This provides the desired properties that the token **406** can be independently validated without requiring a coordinated global update in order to evolve (or revoke) authority ownership.

FIG. **5** shows authorities **402** in a storage cluster **160** executing transactions, using tokens **406**. Each authority **402** has allocated or has been allocated one or more partitions **910** of storage memory in the storage units **152**, e.g. partitions **910** in flash memory **206** and NVRAM **204**. Each authority **402** uses the respective allocated partitions **910**, for writing or reading data. Authorities can be associated with differing amounts of physical storage of the system in some embodiments. For example, one authority **402** could have a larger number of partitions **910** or larger sized partitions **910** in one or more storage units **152** than one or more other authorities **402**.

An operation or command involving user data or metadata could originate at one of the authorities **402**, or could originate external to the storage cluster **160**, for example an I/O request or other command. The operation could be executed by the authority **402** generating or receiving it, or could be forwarded to another authority **402** or even generated by multiple authorities **402** and handed to an authority **402**. Each operation can include one or more transactions. Transaction processing, for an operation or command, is discussed below in general terms, and followed by illustration in the form of an example.

Transaction information, associated with an operation, is provided to one or more authorities **402**, so that each such authority **402** can execute one or more transactions. In some embodiments, the transaction information includes data that enables a restart of the operation. A token **406** is assigned to each transaction and the authority **402** associated with the transaction. This assignment can be done by the authority **402** that receives the transaction information and executes the transaction. In some embodiments, each token **406** is specific to a respective transaction. In other embodiments, the token **406** is specific to the operation and a copy of the token **406** is assigned to or accompanies each transaction for the operation. Each transaction may be executed in accordance with the dependencies described in the transaction information, by the authority **402** that receives the transaction information for that transaction. In some embodiments,

while this is occurring, the system maintains a pre-transaction state of the system, with respect to that operation. Transactions, or portions of transactions, can be processed in parallel, while a consistent view is provided across the system. Multiple commands, each with multiple transactions, can be processed through multiple authorities in parallel. Any inquiry as to system state receives a consistent answer, and any failure during transaction processing leaves the system in the pre-transaction state for consistent recovery from the failure.

The authority 402 that owns the transaction (for example, an authority 402 identified as A_1) communicates to each of the other authorities (for example, authorities 402 identified as A_2 through A_n) to begin their respective portions of the transaction. Then each of the other authorities, A_2 through A_n , in turn, prepare any changes they might have to make for the transaction and pass that information along with a token that the other authorities (A_2 through A_n) generated back to the authority 402 A_1 , that owns the transaction. Authority A_1 that owns the transaction then commits the entire set of changes to NVRAM 204 of the storage unit 152 as collected from each of the other authority (A_2 through A_n) along with the tokens 406 that they provided to the storage devices. The token 406 provides proof to the storage tier that the update to that authority's data is still allowed (protecting against delays in message transmission or authority failovers, etc.). The authority 402 that is committing the entire set of changes to the NVRAM 204 writes the data and/or metadata associated with that transaction, and the token 406 associated with that transaction, into a partition 910 in a storage unit 152 associated with that authority 402. Each of the authorities 402 is capable of receiving transaction information and generating a token 406, executing the transaction and communicating changes for the transaction, and the token 406, to the authority 402 that owns the transaction. Each of the authorities 402 can receive commands, and each of the authorities 402 can act as a transaction authority for any authority receiving commands. In some embodiments, each of the tokens 406 has a validity interval and/or an expiration time, established at the time of creating the token. A time-based token 406 indicates permission to perform the associated transaction within a time span. In transactions involving data writing, the token 406 permits or enables an authority 402 to write the data into a partition that is allocated to that authority 402. If a token 406 expires, and the corresponding transaction has not been performed, the system is maintained at the pre-transaction state. That is, the operation itself is not committed if any transaction of the operation has not been completed by the time a corresponding token 406 associated with that transaction expires. Any of the authorities 402, and particularly the authority 402 requesting or governing the operation with the multiple transactions, can access tokens 406 to see if a token 406 has expired in some embodiments. Any of the authorities 402 can check to see if data has been written to a partition 910. Writing data associated with each of one or more transactions for a command indicates that the command is committed. That is, the command is committed once each transaction of the command is written into an allocated partition of the storage device. It should be appreciated that this feature is useful in determining system state, for replies to inquiries, and is also useful in recovery from failure. In some embodiments, an authority 402 executing a transaction (i.e., a transaction authority 402) logs the transaction, for example by writing a log in NVRAM 204. The log can be persisted, and can be accessed by other authorities 402 to assist in recovery from failure. In an alternate embodiment,

an authority 402 that owns a transaction can write the associated token 406 into the corresponding partition 910 assigned to the authority that owns the transaction. One or more of the above mechanisms and techniques is used in various embodiments, to maintain and determine a consistent system state.

After all of the transactions of an operation that has multiple transactions have completed, with verification available using one of the above-described mechanisms, the system can transition to a post-transaction state. The system considers the operation complete and committed only when all of the transactions associated with a particular operation have been executed, e.g., data written into allocated portions of the storage device, as permitted by a corresponding token, and only then does the system transition to the post-transaction state with respect to that operation or command. In embodiments using optimistic transaction execution that turn out to be based on premises later found not valid at the time of the optimistic execution are revised or returned to an original state.

As an example of transaction processing using tokens 406 in a storage cluster 106, consider an I/O request or command to move a file from one directory to another directory. A storage node 150 receiving the I/O request sends information to another authority 402, for example the authority 402 labeled "A1". This authority 402 A_1 takes ownership of the operation, breaks down the operation, namely the I/O request, into multiple transactions, and sends transaction information to each of the authorities 402 affected by the operation. An authority 402 that has ownership of the first directory receives transaction information to delete the reference to the affected file from the first directory, for the file move in this example. An authority 402 that has ownership of the second directory receives transaction information to add reference to the affected file to the second directory, for the file move. The authorities 402 that are executing the transactions generate tokens 406 associated with the transactions and send the tokens along with transaction information and information about changes they have to make for any transaction back to the authority that owns the transaction. The authority that owns the transaction commits the entire set of changes along with the tokens 406 to the partition 910 in the NVRAM 204. As noted above, each of the tokens 406 may be time bound and permit an authority write access to an allocated partition in a storage device. While the transactions are being executed, the system maintains the pre-transaction state, which shows, in this example that the affected file is under the first directory. Only after all transactions associated with the operation of moving the file from the first directory to the second directory have been executed, does the system show the post-transaction state, which is that the file has been moved from under the first directory to under the second directory. All of these transactions are performed without a global lock, so that other transactions of other commands can be performed in parallel with the transactions associated with moving a file from under one directory to under another directory. Any other transactions or commands will see a consistent state of the system. Any time there is a disruption or a failure, the system can recover to whichever consistent system state is in place. Any of the authorities 402 can read data in partitions 910, authority logs, or tokens 406, and can determine what that consistent system state should be, for purposes of inquiries or recovery. Since this information is available for access by authorities 402, there is no need for broadcasting completion of any of the transactions or the operation in the distributed storage system. Similar examples of other types of opera-

tions or commands, such as writing a file, reading a file, deleting a file, copying a file and so on, with associated transactions, tokens, and system states, are readily devised.

FIG. 6 is a flow diagram of a method for processing transactions associated with a command, which can be practiced in the storage cluster of FIGS. 1-5 and further embodiments thereof. Some or all of the actions in the method can be performed by various processors, such as processors in storage nodes or processors in storage units. In an action 602, partitions in storage devices are allocated to authorities. A command is received at a first authority, in an action 604. A pre-transaction state of the system is maintained, in an action 606. In an action 608, a transaction, relating to the command, is sent from the first authority to a second authority. The transaction is accompanied by or associated with a token which can be generated by the second authority. The token is specific to the transaction and indicates in some embodiments that the authority executing the transaction has permission to perform the transaction and is permitted to write data to an allocated partition of a storage device. In an action 610, data is written into a partition in memory of a storage device, by the second authority in accordance with the transaction, and as permitted by the token.

In a decision action 612, it is determined whether the operation is complete. If all transactions associated with the operation have been executed, the operation is considered complete, and flow proceeds to the action 614. Completion of all of the transactions of an operation can be determined through one of the mechanisms described with reference to FIG. 5, such as all of the data associated with each transaction of the command being written into respective allocated portions of a storage device. If the operation is not complete, flow branches back to the decision action 612 to await completion of all the transactions. In the action 614, the system transitions to a post-transaction state of the system.

It should be appreciated that the methods described herein may be performed with a digital processing system, such as a conventional, general-purpose computer system. Special purpose computers, which are designed or programmed to perform only one function may be used in the alternative. FIG. 7 is an illustration showing an exemplary computing device which may implement the embodiments described herein. The computing device of FIG. 7 may be used to perform embodiments of the functionality for distributed transaction processing in accordance with some embodiments. The computing device includes a central processing unit (CPU) 701, which is coupled through a bus 705 to a memory 703, and mass storage device 707. Mass storage device 707 represents a persistent data storage device such as a floppy disc drive or a fixed disc drive, which may be local or remote in some embodiments. The mass storage device 707 could implement a backup storage, in some embodiments. Memory 703 may include read only memory, random access memory, etc. Applications resident on the computing device may be stored on or accessed via a computer readable medium such as memory 703 or mass storage device 707 in some embodiments. Applications may also be in the form of modulated electronic signals modulated accessed via a network modem or other network interface of the computing device. It should be appreciated that CPU 701 may be embodied in a general-purpose processor, a special purpose processor, or a specially programmed logic device in some embodiments.

Display 711 is in communication with CPU 701, memory 703, and mass storage device 707, through bus 705. Display

711 is configured to display any visualization tools or reports associated with the system described herein. Input/output device 709 is coupled to bus 705 in order to communicate information in command selections to CPU 701. It should be appreciated that data to and from external devices may be communicated through the input/output device 709. CPU 701 can be defined to execute the functionality described herein to enable the functionality described with reference to FIGS. 1-6. The code embodying this functionality may be stored within memory 703 or mass storage device 707 for execution by a processor such as CPU 701 in some embodiments. The operating system on the computing device may be iOS™, MS-WINDOWS™, OS/2™, UNIX™, LINUX™, or other known operating systems. It should be appreciated that the embodiments described herein may also be integrated with a virtualized computing system implemented with physical computing resources. Detailed illustrative embodiments are disclosed herein. However, specific functional details disclosed herein are merely representative for purposes of describing embodiments. Embodiments may, however, be embodied in many alternate forms and should not be construed as limited to only the embodiments set forth herein.

It should be understood that although the terms first, second, etc. may be used herein to describe various steps or calculations, these steps or calculations should not be limited by these terms. These terms are only used to distinguish one step or calculation from another. For example, a first calculation could be termed a second calculation, and, similarly, a second step could be termed a first step, without departing from the scope of this disclosure. As used herein, the term “and/or” and the “/” symbol includes any and all combinations of one or more of the associated listed items.

As used herein, the singular forms “a”, “an” and “the” are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms “comprises”, “comprising”, “includes”, and/or “including”, when used herein, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof. Therefore, the terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting.

It should also be noted that in some alternative implementations, the functions/acts noted may occur out of the order noted in the figures. For example, two figures shown in succession may in fact be executed substantially concurrently or may sometimes be executed in the reverse order, depending upon the functionality/acts involved.

With the above embodiments in mind, it should be understood that the embodiments might employ various computer-implemented operations involving data stored in computer systems. These operations are those requiring physical manipulation of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. Further, the manipulations performed are often referred to in terms, such as producing, identifying, determining, or comparing. Any of the operations described herein that form part of the embodiments are useful machine operations. The embodiments also relate to a device or an apparatus for performing these operations. The apparatus can be specially constructed for the required purpose, or the apparatus can be a general-purpose computer selectively activated or configured by a

computer program stored in the computer. In particular, various general-purpose machines can be used with computer programs written in accordance with the teachings herein, or it may be more convenient to construct a more specialized apparatus to perform the required operations.

A module, an application, a layer, an agent or other method-operable entity could be implemented as hardware, firmware, or a processor executing software, or combinations thereof. It should be appreciated that, where a software-based embodiment is disclosed herein, the software can be embodied in a physical machine such as a controller. For example, a controller could include a first module and a second module. A controller could be configured to perform various actions, e.g., of a method, an application, a layer or an agent.

The embodiments can also be embodied as computer readable code on a non-transitory computer readable medium. The computer readable medium is any data storage device that can store data, which can be thereafter read by a computer system. Examples of the computer readable medium include hard drives, network attached storage (NAS), read-only memory, random-access memory, CD-ROMs, CD-Rs, CD-RWs, magnetic tapes, and other optical and non-optical data storage devices. The computer readable medium can also be distributed over a network coupled computer system so that the computer readable code is stored and executed in a distributed fashion. Embodiments described herein may be practiced with various computer system configurations including hand-held devices, tablets, microprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers and the like. The embodiments can also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a wire-based or wireless network.

Although the method operations were described in a specific order, it should be understood that other operations may be performed in between described operations, described operations may be adjusted so that they occur at slightly different times or the described operations may be distributed in a system which allows the occurrence of the processing operations at various intervals associated with the processing.

In various embodiments, one or more portions of the methods and mechanisms described herein may form part of a cloud-computing environment. In such embodiments, resources may be provided over the Internet as services according to one or more various models. Such models may include Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS). In IaaS, computer infrastructure is delivered as a service. In such a case, the computing equipment is generally owned and operated by the service provider. In the PaaS model, software tools and underlying equipment used by developers to develop software solutions may be provided as a service and hosted by the service provider. SaaS typically includes a service provider licensing software as a service on demand. The service provider may host the software, or may deploy the software to a customer for a given period of time. Numerous combinations of the above models are possible and are contemplated.

Various units, circuits, or other components may be described or claimed as “configured to” perform a task or tasks. In such contexts, the phrase “configured to” is used to connote structure by indicating that the units/circuits/components include structure (e.g., circuitry) that performs the task or tasks during operation. As such, the unit/circuit/

component can be said to be configured to perform the task even when the specified unit/circuit/component is not currently operational (e.g., is not on). The units/circuits/components used with the “configured to” language include hardware—for example, circuits, memory storing program instructions executable to implement the operation, etc. Reciting that a unit/circuit/component is “configured to” perform one or more tasks is expressly intended not to invoke 35 U.S.C. 112, sixth paragraph, for that unit/circuit/component. Additionally, “configured to” can include generic structure (e.g., generic circuitry) that is manipulated by software and/or firmware (e.g., an FPGA or a general-purpose processor executing software) to operate in manner that is capable of performing the task(s) at issue. “Configured to” may also include adapting a manufacturing process (e.g., a semiconductor fabrication facility) to fabricate devices (e.g., integrated circuits) that are adapted to implement or perform one or more tasks.

The foregoing description, for the purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the embodiments and its practical applications, to thereby enable others skilled in the art to best utilize the embodiments and various modifications as may be suited to the particular use contemplated. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

What is claimed is:

1. A method, comprising:
 - receiving a command;
 - generating a transaction responsive to receiving the command, wherein a token associated with the transaction indicates that a data structure having ownership of the transaction has permission to perform the transaction; and
 - writing data into storage memory in a storage node of the storage system that is allocated to the data structure in accordance with the transaction, wherein the storage node includes differing capacities of solid state storage.
2. The method of claim 1, wherein command processing and transaction processing are performed without a global lock on the storage system.
3. The method of claim 1, wherein while each transaction associated with the command is being executed, the system maintains a pre-transaction state.
4. The method of claim 1, wherein each token is time-based to enable performing the associated transaction within a time span.
5. The method of claim 1, further comprising:
 - processing multiple commands, each having multiple transactions, through multiple data structures in parallel, wherein the writing data associated with each of the multiple transactions for one of the multiple commands indicates the one of the multiple commands is committed.
6. The method of claim 1, further comprising:
 - avoiding broadcasting completion of the transaction throughout the storage system.
7. The method of claim 1, further comprising:
 - committing the command once each transaction of the command is written into an allocated partition.

19

8. A tangible, non-transitory, computer-readable media having instructions thereupon which, when executed by a processor, cause the processor to perform a method comprising:

receiving command;

generating a transaction responsive to receiving the command, wherein a token associated with the transaction indicates that a data structure having ownership of the transaction has permission to perform the transaction; and

writing data into storage memory in a storage node of the storage system that is allocated to the data structure in accordance with the transaction, wherein the storage node includes differing capacities of solid state storage.

9. The computer-readable media of claim **8**, wherein command processing and transaction processing are performed without a global lock on the storage system.

10. The computer-readable media of claim **8**, wherein while each transaction associated with the command is being executed, the system maintains a pre-transaction state.

11. The computer-readable media of claim **8**, wherein each token is time-based to enable performing the associated transaction within a time span.

12. The computer-readable media of claim **8**, wherein the method further comprises:

processing multiple commands, each having multiple transactions, through multiple data structures in parallel, wherein the writing data associated with each of the multiple transactions for one of the multiple commands indicates the one of the multiple commands is committed.

13. The computer-readable media of claim **8**, wherein broadcasting completion of the transaction throughout the storage system is avoided.

14. The computer-readable media of claim **8**, wherein the method further comprises:

committing the command once each transaction of the command is written into an allocated partition.

20

15. A storage system comprising:

a memory; and

a processing unit having a processor configured to execute actions, the actions comprising:

receiving command;

generating a transaction responsive to receiving the command, wherein a token associated with the transaction indicates that a data structure having ownership of the transaction has permission to perform the transaction; and

writing data into storage memory in a storage node of the storage system that is allocated to the data structure in accordance with the transaction, wherein the storage node includes differing capacities of solid state storage.

16. The storage system of claim **15**, wherein command processing and transaction processing are performed without a global lock on the storage system.

17. The storage system of claim **15**, wherein while each transaction associated with the command is being executed, the system maintains a pre-transaction state.

18. The storage system of claim **15**, wherein each token is time-based to enable performing the associated transaction within a time span.

19. The storage system of claim **15**, wherein the actions further comprise:

processing multiple commands, each having multiple transactions, through multiple data structures in parallel, wherein the writing data associated with each of the multiple transactions for one of the multiple commands indicates the one of the multiple commands is committed.

20. The storage system of claim **15**, wherein broadcasting completion of the transaction throughout the storage system is avoided.

* * * * *